Costly decisions and sequential bargaining

James Costain
Banco de España

December 2014

Abstract

This paper models a near-rational agent who chooses from a set of feasible alternatives, subject to a cost function for precise decision-making. Unlike previous papers in the "control costs" tradition, here the cost of decisions is explicitly interpreted in terms of time. That is, by choosing more slowly, the decision-maker can achieve greater accuracy. Moreover, the *timing* of the choice is itself also treated as a costly decision.

A tradeoff between the precision and the speed of choice becomes especially interesting in a strategic situation, where each decision maker must react to the choices of others. Here, the model of costly choice is applied to a sequential bargaining game. The game closely resembles that of Perry and Reny (1993), in which making an offer, or reacting to an offer, requires a positive amount of time. But whereas Perry and Reny treat the decision time as an exogenous fixed cost, here we allow the decision-maker to vary precision by choosing more or less quickly.

Numerical simulations of bargaining equilibria closely resemble those of the Binmore, Rubinstein, and Wolinsky (1983) framework, except that the time to reach agreement is nonzero. In contrast to the model of Perry and Reny, we find that rejecting an offer and proposing an alternative are not equivalent, and that equilibrium is unique when the space of possible offers is continuous.

Keywords: Bargaining, control costs, logit equilibrium, near rationality, sticky wages.
JEL Codes: C72, C78, D81

1 Introduction[1]

Frictions are essential in macroeconomic modeling. Empirically-oriented DSGE models, following Christiano, Eichenbaum, and Evans (2005), Smets and Wouters (2003), and Gertler, Sala, and Trigari (2008), often feature sticky prices, sticky wages, investment adjustment costs, consumption habits, and search and matching in the labor market, among other frictions. Similarly, errors are essential in game theory. Theoretical models often fit experimental data better under equilibrium concepts that incorporate errors in choice, such as quantal response equilibrium (McKelvey and Palfrey 1995, 1998), or its special case, logit equilibrium. In addition, equilibrium concepts such as trembling hand equilibrium (Selten 1975), quantal response equilibrium,

[1]Discussions with Galo Nuño and Filip Matejka were very helpful for getting this project started; the paper builds on previous work with Anton Nakov. I have also received helpful comments from Henrique Basso, Espen Moen, Pascal Michaillat, Plamen Nenov, and Ernesto Villanueva, and from seminar participants at the Banco de España, the Univ. of Murcia, UC Santa Cruz, Univ. Carlos III, and CEF2014. Views expressed here are those of the author and do not necessarily coincide with those of the Bank of Spain or the Eurosystem.

and control cost equilibrium (Stahl 1990; Van Damme, 1991, Chapter 4; Mattsson and Weibull 2002) have been useful for resolving some behavioral puzzles associated with fully rational Nash equilibria, and as robustness criteria for selecting between multiple equilibria that occur under full rationality (Goeree and Holt 1999, 2001; Anderson, Goeree, and Holt 2002).

Control cost equilibria are based on the assumption that errors occur because decisions are costly. In particular, a decision is conceived as a random variable distributed over a set of possible actions, and the cost of the decision is assumed to be an increasing function of the precision of that random variable. The player maximizes the payoffs that would obtain in the fully rational game, net of decision costs. Typically, a player optimally spreads probability across many possible actions (thus committing "errors") rather than concentrating all probability on a single action, because the latter is excessively costly. Thus, players are sufficiently sophisticated to consider the costs and limitations of their own choice process when they make decisions, which is an appealing property. Note, however, that the costs of choice are likely to include *time used up in the decision.* This suggests that the time taken up by choice should not only be subtracted out of the player's net payoffs— it should also have strategic implications. That is, time used up on one decision is time that cannot be devoted to other decisions, and may represent an opportunity for other players to take actions of their own. Therefore, the time used up in decisions *should be reflected in the extensive form of the game.*

Accordingly, this paper explores the implications of control cost equilibrium when we take seriously the role of time in decision-making. The model, developed in Section 2, studies a decision-maker (DM) who can choose quickly or slowly, and can make a more accurate decision by choosing more slowly.[2] As in previous papers on control costs, "choosing" means allocating probability across a set of feasible actions.[3] Unlike previous papers, the DM is also assumed to control the *arrival rate* of the decision. Holding fixed other uses of time, a slower arrival rate implies more time dedicated to the decision, and this, by assumption, permits a more precise allocation of probabilities across the action set. While a variety of statistics could serve as measures of precision, this paper focuses on a special case in which precision is measured by relative entropy. This proves analytically convenient, because the probability of each feasible action can then be written in the form of a logit (as in Mattsson and Weibull, 2002). More precisely, we will see that measuring precision by relative entropy helps us obtain an interior solution, and guarantees some convenient invariance properties that help us understand the implied choice behavior.

In addition to the time-consuming, error-prone choice across feasible actions, the model treats the *timing* of the choice as a time-consuming, error-prone decision too. Under the functional forms assumed in the paper, the timing of the choice is determined by a weighted binary logit (as in Woodford, 2008). Considering errors on both margins— the choice itself, and the timing of the choice— actually simplifies the analysis, because it helps rule out corner solutions. While the model can have several types of corner solutions in the limiting case where timing is perfectly rational (see Section 2.1), it has a well-behaved interior solution when timing is error prone (Section 2.4). Interestingly, the model implies a relationship between the accuracy of decision-making across the two margins (the choice itself, and its timing), as well as a relation between the accuracy of decision-making and the value of time. These implications of the model may be empirically testable, especially in the laboratory.

As an application, Section 3 builds the model of costly decisions into a game where two

[2]Rubinstein (2013) presents experimental evidence that slower decisions are more accurate.

[3]For a discussion of decisions as random variables, see Machina (1985) or Chapter 2 of Anderson *et al.* (1992).

players bargain to split a pie. Time-consuming, error-prone choice is assumed both at the stage of offering a share to the other player, and at the stage of accepting or rejecting the other player's proposal. The game closely resembles that of Perry and Reny (1993), in which two players make offers to split a pie, and making an offer, or reacting to an offer, requires a nonnegative amount of time. But whereas Perry and Reny treat the decision time as an exogenous fixed cost, and assume that all decisions are optimal, here the decision-maker can vary precision by choosing more or less quickly. Also, while Perry and Reny equate rejecting an offer with making an alternative offer, the present model implies that the former is simpler than the latter. Therefore the average equilibrium time to agreement is longer in Sec. 3.1, where we simulate a bargaining game in which rejection requires a player to propose an alternative, than it is under the protocol studied in Sec. 3.2, where players can just say "no".

Sections 3.3-3.5 further explore the structure of bargaining equilibria by numerical simulation. The equilibria are very similar to those of the Binmore, Rubinstein, and Wolinsky (1983) framework, except that the time to reach agreement is nonzero. The simulations explore the comparative statics of the game, including the case where splitting the pie is less valuable than the option of continuing without agreement. Finally, we explore uniqueness of equilibrium. While the game of Perry and Reny (1993) has multiple equilibria in which players receive different bargaining shares, in our game we find a unique equilibrium as long as the space of possible offers is continuous. Section 4 concludes.

1.1 Related literature

This paper is closely related to previous work on sequential bargaining under complete information, including Rubinstein (1982), and Binmore, Rubinstein, and Wolinsky (1983). Wolinsky (1987) and Hall and Milgrom (2005) have emphasized the implications of bargaining theory for wages in matching models. Perry and Reny (1993) studied bargaining when making offers, or responding to offers, requires a fixed, nonnegative quantity of time. Merlo and Wilson (1995, 1998) and Merlo and Tang (2010) study bargaining games where the pie may evolve randomly over time.

However, the motivation for this paper comes more directly from the author's previous work on price stickiness. Costain and Nakov (2014) showed that a control cost approach is fruitful for modeling microdata on intermittent retail price adjustment. They showed that two types of errors were relevant for the empirical success of their model. Errors in *which price to set*, conditional on adjustment, help explain a number of empirical puzzles on retail prices; errors in the *timing of adjustment* help explain the macroeconomic finding of significant monetary nonneutrality. They argued that their model of intermittent, error-prone adjustment was potentially applicable to many contexts other than retail prices; the present paper explores how the framework can be extended to analyze play in extensive-form games.

Another motivation for the present paper is that it may help unify the analysis of several margins present in labor market models. Cheremukhin, Restrepo-Echevarria, and Tutino (2012) have shown how a matching function can be derived from a model of costly choice over a set of partners (who are themselves, likewise, engaged in costly choice). The present paper uses a cost function very similar to that of Cheremukhin *et al.* to describe another labor market phenomenon, namely, wage bargaining. Since the time to reach agreement is strictly positive (in contrast to Rubinstein (1982) and many related models) the present model could also be applied to duration data on wage negotiations and/or strikes. After a bargain is accepted, the same model of error-prone timing could be applied to each partner's option to reopen the negotiations, which is a natural way of modeling wage stickiness. Likewise, this decision model could be applied to each partner's separation decision. Indeed, wage stickiness and separation could be treated jointly in this framework, as Barro (1977) advocated.

Thus, control cost equilibrium offers a single microfoundation for a number of frictions that are typically viewed as distinct, including nominal rigidities and matching frictions. Applying a single model of frictions to all margins could make them easier to calibrate and compare, and offers empirical implications about how the degree of errors varies across margins over time. Control cost models were initially developed as a more structured alternative to trembling hand equilibrium, imposing the property that more costly errors should be less likely (van Damme, 1991, Chapter 4). Starting with Stahl (1990), many papers have shown independently how an entropy-related cost function can microfound logit decision rules (Marsili, 1999; Mattsson and Weibull, 2002; Bono and Wolpert, 2009; Matejka and MacKay, 2011). Similar cost functions have been used to model other limitations on rationality, as in the model uncertainty approach of Hansen and Sargent (2007). Control cost models are also influential in the engineering and machine learning literature; see for example Todorov (2009) and Theodorou, Dvijotham, and Todorov (2013). In the machine learning context, as in the reinforcement learning literature in economics (see Baron *et al.*, 2002), control costs are applied to backward-looking behavior, whereas in the present paper they are applied to forward-looking behavior.

The rational inattention model of Sims (1998, 2003) is also a general friction applicable to many different types of decisions, and it is quite similar to a control cost approach. Indeed, the papers of Woodford (2008) and Cheremukhin *et al.* (2012) both study rational inattention models. The main difference between a rational inattention model and a control cost model is that the former places a constraint on information flow (measured in terms of entropy), while the latter places a constraint on the precision of the decision (entropy is then one possible functional form for measuring precision). In other words, the frameworks address two different "stages" of the decision process: the initial stage of obtaining information necessary for the decision, and the final stage of actually making a choice conditional on that information. In reality, both these stages are likely to be costly, and both have the same primary implication: an imperfect correlation between an agent's true state and its action. But there is an important technical advantage to modeling the second stage, rather than the first: the rational inattention approach implies a much higher-dimensional model, since the decision-maker's state variable is his prior (which in most contexts is a high-dimensional object). Moreover, the present approach is particularly well-suited to describing control variables that are adjusted intermittently, making this model applicable in a variety of interesting economic contexts, such as nominal stickiness, formation and dissolution of relationships, and portfolio adjustment, to name just a few.

2 Choosing when to solve a problem: a "control cost" approach

2.1 One costly decision

The paper begins by analyzing a single costly decision, which will subsequently form the building blocks of a sequential bargaining game. We regard a decision as a random variable distributed across a set of possible alternatives. The cost of making a choice is that it requires time. The

Figure 1: A decision that requires time.

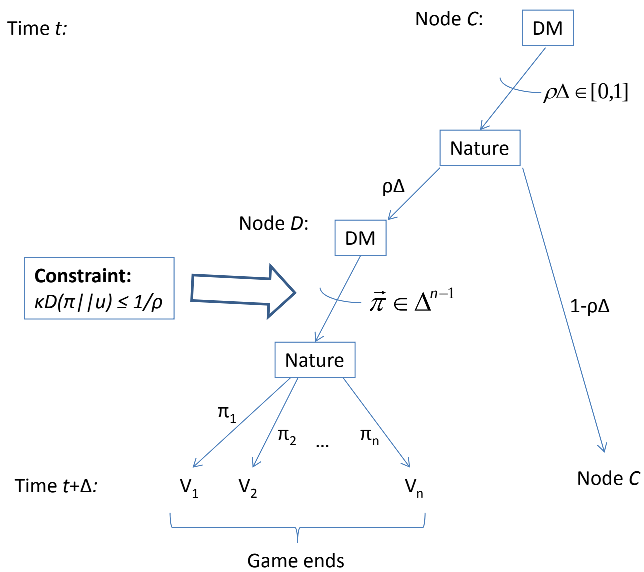

Note: Decision maker (DM) chooses arrival rate ρ of decision. Nature lets decision arrive with probability $\rho\Delta$. A slower decision implies higher precision in the allocation of probabilities π_i across alternatives i with values V_i, satisfying the constraint $\kappa\mathcal{D}(\vec{\pi}||\vec{u}) \leq \rho^{-1}$.

decision-maker faces a tradeoff: if she chooses more quickly, her decision will be less precise.[4] Precision is measured in terms of relative entropy (also known as Kullback-Leibler divergence). A perfectly random (that is, uniform) decision is costless; the time cost of any decision that deviates from uniformity is proportional to the relative entropy between that decision and a uniform decision. That is, for a given decision $\vec{\pi}$, which we write as a vector of probabilities associated with all the possible alternatives that could be chosen, the time cost is $\kappa\mathcal{D}(\vec{\pi}||\vec{u})$, where $\vec{u}$ is a uniform distribution over the same alternatives.[5]

It would be inelegant, unrealistic, and inconvenient to assume that the time required for a given decision is exactly known. So rather than assuming that the DM chooses the decision time directly, we assume that the choice variable is the arrival rate of the decision, ρ. The expected duration of the decision is (approximately) $1/\rho$; a slower arrival rate implies more time spent making the decision, and we assume this permits a more precise allocation of probability across the various alternatives under consideration.

The structure of this choice process is illustrated by the game tree in Figure 1, which shows one time step in the decision process, occurring at time t. The continuous-time environment is approximated by a sequence of steps of duration Δ. The decision maker (DM) chooses the

[4]The model developed in this subsection is analogous to the "PPS-control" specification considered in Costain and Nakov (2014): it allows for errors in a choice across a set of options, but implicitly assumes that the *timing* of the choice is perfectly optimal.

[5]The notation used here to denote the relative entropy $\mathcal{D}(\vec{p}||\vec{q})$ between distributions $\vec{p}$ and $\vec{q}$ is standard. When the two probability vectors have length n, the Kullback-Leibler divergence is defined as $\mathcal{D}(\vec{p}||\vec{q}) \equiv \sum_{i=1} p_i \ln(p_i/q_i)$. For a discussion, see Cover and Thomas (2006), Chapter 2.3.

arrival rate ρ of the decision at node C ("calculating"); the curve lying below the node indicates the assumption that the hazard is chosen from a continuum of possible values, $\rho \in [0, \Delta^{-1}]$. Since the actual time to complete the decision is stochastic, we thereafter see a choice by Nature between finishing the decision with probability $\rho\Delta$, and continuing the process with probability $1 - \rho\Delta$, in which case the game returns to a node of type C at time $t + \Delta$.

When Nature allows the decision to be completed, the game moves to node D ("deciding"), at which point the DM allocates probabilities π_i across the alternatives $i \in \{1, 2, ...n\}$ with values V_i. The precision of the decision $\vec{\pi} \equiv (\pi_1, \pi_2, \ldots, \pi_n)' \in \Delta^{n-1}$ is limited by the rate at which the choice arrived (since $\vec{\pi}$ is a probability vector of length n, it is chosen from the simplex of order $n - 1$, which we write as Δ^{n-1}). Concretely, we assume

$$\frac{1}{\rho} \;\geq\; \kappa \mathcal{D}(\vec{\pi}||\vec{u}) \;\equiv\; \kappa \left(\sum_{i=1}^{n} \pi_i \ln \pi_i + \ln n \right). \tag{1}$$

That is, a slower decision can have greater precision, where precision is measured by the Kullback-Leibler divergence between the decision $\vec{\pi}$ and a uniform distribution $\vec{u}$.[6]

The model is written assuming discreteness in two dimensions: a choice set of n discrete options, and discrete time steps of length Δ. However, we will focus on deriving solutions that are independent of this discreteness. In particular, it is helpful to measure precision by *relative entropy* rather than by entropy *per se*. Entropy does not have a finite limit as it is calculated on finer and finer grids; in contrast, relative entropy is invariant as the density of grid points increases (see Cover and Thomas, 2006, Chapter 8). Therefore, the discreteness assumed in this section is only a notational convention, and indeed, when we apply the decision model inside a bargaining game in Section 3, we will allow for a continuous support of possible offers.

The delay required for more precise decisions could have a variety of costs. As in Rubinstein (1982), this paper first considers the case in which the only cost of delay is pure time discounting, at rate δ. Now, let W_t denote the value of the problem at node C at time t.[7] Approximating the continuous-time problem by a discrete-time problem with time steps of length Δ, W_t can be derived from the following Bellman equation:

$$\begin{aligned} W_t \quad &= \max_{\rho,\, \{\pi_i\}_{i=1}^n} \; (1 - \delta\Delta) \left[\rho\Delta \sum_{i=1}^{n} \pi_i V_i + (1 - \rho\Delta) W_{t+\Delta} \right] \\ \text{subject to:} \quad & \kappa \left(\sum_{i=1}^{n} \pi_i \ln \pi_i + \ln n \right) \;\leq\; \frac{1}{\rho} \\ \text{and} \quad & \sum_{i=1}^{n} \pi_i = 1. \end{aligned} \tag{2}$$

[6] Uniformity of the benchmark distribution $\vec{u}$ is just a simplifying assumption; it is not essential for our model. The model could be developed assuming some other benchmark distribution, just as it could be developed assuming some measure of precision other than relative entropy.

[7] We introduce a time subscript on W because when we apply this model inside a larger dynamic game later in the paper, the value of the problem may be changing over time. We will solve the larger game by backwards induction, so the time $t + \Delta$ value $W_{t+\Delta}$ will be taken as given when solving the time t problem, which makes it important to distinguish between W_t and $W_{t+\Delta}$. Other time subscripts will be suppressed, to lighten the notation.

We form the following Lagrangian:

$$\mathcal{L} \;=\; (1-\delta\Delta)\left[\rho\Delta\sum_{i=1}^{n}\pi_i V_i + (1-\rho\Delta)W_{t+\Delta}\right] + \lambda_\rho\left[1-\rho\kappa\left(\sum_{i=1}^{n}\pi_i\ln\pi_i + \ln n\right)\right] + \lambda_\pi\left(1-\sum_{i=1}^{n}\pi_i\right).$$

The first-order conditions of the problem are:

$$\begin{aligned}
\rho\Delta(1-\delta\Delta)V_i - \lambda_\rho\rho\kappa(1+\ln\pi_i) - \lambda_\pi &= 0 && (3)\\
\Delta(1-\delta\Delta)\left(\sum_i \pi_i V_i - W_{t+\Delta}\right) - \lambda_\rho\kappa\mathcal{D}(\vec{\pi}||\vec{u}) &= 0 && (4)\\
1-\rho\kappa\mathcal{D}(\vec{\pi}||\vec{u}) &= 0 && (5)\\
1-\sum_i \pi_i &= 0. && (6)
\end{aligned}$$

Rearranging, the necessary condition for π_i can be rewritten as

$$\pi_i \;=\; \exp\left(\beta V_i - \frac{\beta\mu}{\rho(1-\delta\Delta)} - 1\right), \tag{7}$$

where

$$\beta \;=\; \frac{(1-\delta\Delta)\Delta}{\kappa\lambda_\rho}\,. \tag{8}$$

Equation (7) shows that π_i is proportional to the exponential of the value of alternative i, scaled by a coefficient β which indicates the degree of rationality with which the decision is taken. A first-order condition of this form gives rise to a logit variable. Since the probabilities π_i must sum to one, (7) implies

$$\pi_i \;=\; \frac{\exp(\beta V_i)}{\sum_j \exp(\beta V_j)}, \tag{9}$$

that is, the probabilities of the various alternatives i are given by a multinomial logit.

It is helpful now to introduce the notation $E^\pi x \equiv \sum_{i=1}^n \pi_i x_i$ to represent the expectation of a random variable x with distribution $\vec{\pi}$. Letting $\vec{u}$ represent a uniform distribution, we have $E^u x = n^{-1}\sum_{i=1}^n x_i$. Note that the optimal distribution (9) implies

$$\pi_i\ln\pi_i \;=\; \beta\pi_i V_i - \pi_i\ln\left(nE^u\exp(\beta V)\right) \tag{10}$$

and therefore

$$\mathcal{D}(\vec{\pi}||\vec{u}) = \sum_{i=1}^{n}\pi_i\ln\pi_i + \ln n = \beta E^\pi V - \ln\left(nE^u\exp(\beta V)\right) + \ln n = \beta E^\pi V - \ln\left(E^u\exp(\beta V)\right) \;\equiv\; K(\beta). \tag{11}$$

Here $E^\pi V$ is the expected value of finishing the decision, which is

$$E^\pi V \;=\; \sum_i \pi_i V_i \;=\; \frac{E^u V\exp(\beta V)}{E^u\exp(\beta V)}. \tag{12}$$

The function $K(\beta)$ represents the cost of achieving precision level β when probabilities are allocated optimally among the alternatives, according to rule (9).

The precision variable β is inversely related to the multiplier λ_ρ. Therefore the first-order condition (4) for the arrival rate ρ can also be written in terms of β, as follows:

$$\beta\left(E^\pi V - W_{t+\Delta}\right) \;=\; \mathcal{D}(\vec{\pi}||\vec{u}). \tag{13}$$

But now we have two different equations for the cost measure $\mathcal{D}$. On one hand, (11) expresses the time cost when the probabilities π_i are allocated optimally across alternatives i. On the other hand, (13) uses the first-order condition for ρ to express an optimal tradeoff between the solution cost $\mathcal{D}$ and the benefit of arriving at the solution, which is $E^\pi V - W_{t+\Delta}$. Combining the two equations, both $\mathcal{D}$ and $\beta E^\pi V$ cancel, and we are left with one simple equation that determines β, given the values $\vec{V}$ of the alternatives and the continuation value $W_{t+\Delta}$:

$$\beta W_{t+\Delta} \;=\; \ln(E^u \exp(\beta V)), \tag{14}$$

or equivalently,[8]

$$E^u \exp(\beta(V - W_{t+\Delta})) \;=\; 1. \tag{15}$$

The function $g(\beta,\vec{V}) \equiv \ln(E^u \exp(\beta V))$ that appears on the right-hand side of (14) is called the *cumulant generating function* of the random variable V under a uniform distribution. (It is the log of the moment generating function.) To clarify the behavior of the equation (14) that determines β, it is helpful to state some properties of the functions g and K.

Lemma 1 *Let $g(\beta,\vec{V}) \equiv \ln(E^u \exp(\beta V)) \equiv \ln\left(\frac{1}{n}\sum_{i=1}^n \exp(\beta V_i)\right)$, where n is finite and all V_i are finite. Then:*

(a.) $g(0,\vec{V}) = 0$.

(b.) $\frac{\partial g}{\partial \beta}(\beta,\vec{V}) = \frac{E^u V \exp(\beta V)}{E^u \exp(\beta V)} = E^\pi V$. *In particular,* $\frac{\partial g}{\partial \beta}(0,\vec{V}) = E^u V$.

(c.) $\frac{\partial^2 g}{\partial \beta^2}(\beta,\vec{V}) = \frac{E^u V^2 \exp(\beta V)}{E^u \exp(\beta V)} - \left(\frac{E^u V \exp(\beta V)}{E^u \exp(\beta V)}\right)^2 = E^\pi V^2 - (E^\pi V)^2$. *Hence for all $\beta \geq 0$, we have $\frac{\partial^2 g}{\partial \beta^2}(\beta,\vec{V}) \geq 0$, with equality if and only if $V_i = E^u V$ for all i.*

(d.) *For all $\beta > 0$, $g(\beta,\vec{V}) \geq \beta E^u V$, with equality if and only if $V_i = E^u V$ for all i.*

(e.) *Let $K(\beta)$ be given by (11). Then $K'(\beta) \geq 0$ for all $\beta > 0$ with equality if and only if $V_i = E^u V$ for all i. Also, $K'(0) = 0$, and $\lim_{\beta\to\infty} K'(\beta) = 0$.*

(f.) *The cost of achieving perfect precision is $\lim_{\beta\to\infty} K(\beta) = \ln n$.*

Points (b) and (c) follow directly by differentiating and then simplifying using the logit probability formula (9). The derivative in (c) is nonnegative since it can be interpreted as a variance: it is the variance of the payoff V when choices are distributed according to the logit probabilities (9). Point (d) follows from the fact that $\exp(x)$ is a convex function. Therefore $E^u \exp(\beta V) \geq \exp(\beta E^u V)$, with equality only if all values V_i are equal, so that $V_i = E^u V$ for all i. Therefore if there is any difference across the values V_i, we have $\ln(E^u \exp(\beta V)) > \ln(\exp(\beta E^u V)) = \beta E^u V$. Part (e) follows by noting that $K(\beta) = \beta\frac{\partial g}{\partial \beta} - g$, and therefore $K'(\beta) = \frac{\partial g}{\partial \beta} + \beta\frac{\partial^2 g}{\partial \beta^2} - \frac{\partial g}{\partial \beta} = \beta\frac{\partial^2 g}{\partial \beta^2}$. Then the nonnegative slope of K follows from (c), and we also obtain $K'(0) = 0$. Point (f)– proved informally, need to write down the details. Since K has a finite limit, we also note that

[8] Equations (14) and (15) are mathematically equivalent. The former is graphed in Figure 2, but the latter may be easier to solve numerically, because subtracting $W_{t+\Delta}$ from the choice values V_i before exponentiating may help avoid numerical overflow.

Figure 2: Optimal choice of rationality β under backwards induction.

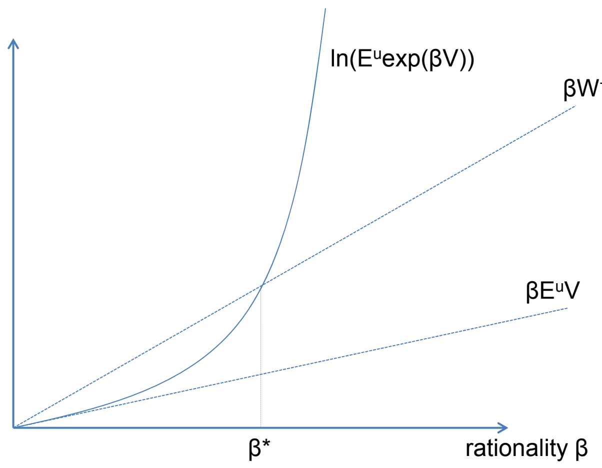

Note: The figure illustrates the interior solution stated in Prop. 1(b). Curve $g(\beta, \vec{V}) = \ln(E^u \exp(\beta V))$ is tangent to $\beta E^u V$ at $\beta = 0$, and has limiting slope $V^* \equiv \max_i V_i$ as $\beta \to \infty$ (the graph assumes $V^* > W^+$).

If instead $V^* \leq W^+$, then $g(\beta, V)$ lies everywhere below βW^+, and postponement is optimal; see Prop. 1(a). If instead $E^u V \geq W^+$, $g(\beta, V)$ lies everywhere above βW^+, and the corner $\beta = 0$ is optimal; see Prop. 1(c).

$\lim_{\beta \to \infty} K'(\beta) = 0$. Thus, as long as there is any variation in the values V_i, the function $K(\beta)$ is S-shaped: it has zero slope at $\beta = 0$ and as $\beta \to \infty$, and strictly positive slope for all $\beta \in (0, \infty)$.

Figure 2 uses these properties of g to show how to solve for the optimal precision, β^*. We first consider the solution of problem (2) when $W_{t+\Delta}$ is taken as a given parameter, as would be the case if we were solving the model by backwards induction. By Lemma 1(b), the curve $g(\beta, \vec{V})$, plotted as a function of β, has slope $E^u V$ at the origin. In the limit as $\beta \to \infty$, the slope of $g(\beta, \vec{V})$ converges to $V^* \equiv \max_i V_i$. Therefore, since g is convex, there exists exactly one positive β^* that solves (14) as long as $E^u V < W_{t+\Delta} < V^*$. This β^* solves (2), as long as the period length Δ is sufficiently short.

Proposition 1 states this conclusion formally, and also describes the corner solutions that arise when the period length is sufficiently long, or when $W_{t+\Delta}$ lies outside of the bounds consistent with an interior solution.[9]

Proposition 1 *Consider problem (2) when $W_{t+\Delta} \equiv W^+$ is a given parameter.*

(a.) Postponement: *Suppose $V^* \equiv \max_i V_i \leq W^+$. Then it is optimal to postpone solving, setting the arrival rate $\rho_t^* = 0$, achieving the value $W_t^* = (1 - \delta\Delta)W^+$.*

(b.) Interior solution: *Suppose $V^* \equiv \max_i V_i > W^+ > E^u V$, and $\frac{\Delta}{\kappa} < \ln n$. Then there exists a unique positive β^* that satisfies (14). Let $\mathcal{D}^*$ and ρ^* be the entropy and arrival rate*

[9]Besides the three solution classes described in Prop. 1, when simulating the model on a computer we must also consider the possibility of numerical overflow solutions. Even if the true solution is theoretically an interior solution β^*, the calculations will overflow on a computer if $\sum_i \exp(\beta^* V_i) > \infty^{NUM}$, where ∞^{NUM} is the largest real number representable on the computer. A similar caveat applies to an immediate solution β^*_Δ. Treatment of these cases will be discussed in an appendix.

given by (11) and (16) when $\beta = \beta^*$. If $\rho^* < \Delta^{-1}$, then β^* solves (2), achieving the value

$$W_t^* = (1 - \delta\Delta)\left[\frac{\Delta\kappa}{\beta^*} + W^+\right].$$

(c.) Immediate solution: *Suppose $E^u V \geq W^+$, or that the conditions of part (b) are satisfied, but the implied arrival probability in the first time step exceeds one: $\rho^*\Delta \geq 1$. Then the optimal arrival rate in problem (2) is $\rho^*_\Delta = \Delta^{-1}$, resulting in a precision level β^*_Δ that solves (17). Alternatively, suppose $\frac{\Delta}{\kappa} \geq \ln n$. Then (2) is solved with infinite precision and arrival rate $\rho^*_\Delta = \Delta^{-1}$.*

The main point is that if $V^* > W_{t+\Delta}$ and the time step is sufficiently short, then an interior solution applies, and precision is given by the unique positive β^* that solves (14). Knowing β^*, we can solve for the other endogenous quantities in the model: π_i and $E^\pi V$ are given by (9) and (12), the precision measure $\mathcal{D}$ is given by (11), and the arrival rate of the decision is

$$\rho \;=\; \frac{1}{\kappa\mathcal{D}(\vec{\pi}||\vec{u})}\,. \tag{16}$$

However, if the value V^* of the best possible option is less than the continuation value $W_{t+\Delta}$, then the decision should simply be postponed. If instead $V^* > W_{t+\Delta}$ and a precision level greater than or equal to β^* can be achieved within a single time step Δ, it is optimal to do so. The precision β^*_Δ of this corner solution can be backed out from the information constraint; if it is finite, it is the unique solution to

$$\beta\frac{E^u V\exp(\beta V)}{E^u\exp(\beta V)} - \ln(E^u\exp(\beta V)) \;=\; \frac{\Delta}{\kappa}. \tag{17}$$

The arrival probability is then $\rho\Delta = 1$ in the first time step. Note, however, that these "immediate" solutions only occur in discrete time; as $\Delta \to 0$ this solution class is eliminated, unless the problem is completely trivial, with $\min_i V_i = \max_i V_i$.

Considering Prop. 1 and Fig. 2, we can draw some comparative statics conclusions about the model's the behavior *conditional* on the continuation value W^+. Fig. 2 shows that conditional on W^+, the solution β is *unaffected* by the information cost parameter κ. Likewise, conditional on W^+, the probabilities π, and $E^\pi V$ and $K(\beta)$ are unaffected by κ. Thus, a rise in κ only affects the arrival rate $\rho = \frac{1}{\kappa K(\beta)}$, slowing down the solution of the problem.

Likewise, we can see from the diagram that an increase in W^+ decreases β^*. Such a change therefore decreases $E^\pi V$ and $K(\beta)$, increasing the arrival rate ρ. Similarly, an increase in the distribution of V in the sense of first-order stochastic dominance, or second-order stochastic dominance (a mean-preserving spread) will increase β^*, since either of these changes will raise the curve $g(\beta) \equiv \ln(E^u\exp(\beta)V)$. This implies that $E^\pi V$ and $K(\beta)$ both increase, while ρ decreases.

The effects of a mean-preserving spread in the distribution of V might seem counterintuitive, since a spread makes it important to choose the right option. But the logit formula (9) shows that a mean-preserving spread of V makes choosing the right option more likely, *even holding β fixed.* That is, implicitly, our model supposes that if the payoff difference between two options is increased, then those options are more easily objectively distinguishable. Thus, a mean-preserving spread of V increases $E^\pi V$ at any $\beta > 0$, and likewise the value of the problem,

$W = \frac{1-\delta\Delta}{\beta\delta\kappa}$, also increases. But since β decreases, the information content $K(\beta)$ of the decision decreases, and thus the arrival rate $\rho = \frac{1}{\kappa\mathcal{D}}$ is increased.

Thus far we have solved (2) at time t, taking as given the continuation value $W_{t+\Delta}$. But it is also interesting to consider an infinite-horizon version of the problem, with time-invariant conditions; in particular, this would require a time-invariant value V_i for each option i. We would then like to calculate the stationary value of the as-yet-unsolved decision problem, which we will denote $W_t = W_{t+\Delta} \equiv W_{ss}$. In this "steady state", we can rearrange the Bellman equation to obtain

$$W_{ss} = (1-\delta\Delta)\frac{\rho}{\delta}(E^{\pi}V - W_{ss}). \tag{18}$$

But using the first-order condition (4) for ρ, we have

$$\beta\left(E^{\pi}V - W_{ss}\right) \;=\; \mathcal{D}(\vec{\pi}||\vec{u}) \;=\; \frac{1}{\rho\kappa}. \tag{19}$$

so by rearranging we find that the steady-state problem value is

$$W_{ss} \;=\; \frac{1-\delta\Delta}{\delta\kappa\beta}. \tag{20}$$

As before, the optimal β must solve $\beta W_{ss} = \ln(E^u \exp(\beta V))$, but now for a steady state problem this simplifies to

$$\ln(E^u \exp(\beta V)) \;=\; \frac{1-\delta\Delta}{\delta\kappa}. \tag{21}$$

The solution of equation (21) is illustrated in Figure 3. Given the solution steady-state interior solution β_{ss}, we can then calculate the entropy $\mathcal{D}_{ss}$, the arrival rate ρ_{ss}, and the value W_{ss} from (11), (16), and (20).

But besides interior solutions, we must again consider some possible corners. It is always feasible to set $\rho = 0$, so that the problem is never solved; this attains the steady-state value $W_{ss} = 0$. Alternatively, the highest possible arrival rate is $\rho = \Delta^{-1}$, so that the decision arrives in exactly one discrete time step of length Δ. We can summarize the range of validity of these different solutions as follows.[10]

Proposition 2 *Consider a time-invariant, infinite horizon version of problem (2). Call the value $W_{ss} \equiv W_t$ for any t.*

(a.) Postponement: *Suppose $V^* \equiv \max_i V_i \leq 0$. Then it is optimal to postpone the solution forever, setting the arrival rate $\rho_{ss} = 0$, achieving the value $W_{ss} = 0$.*

(b.) Interior solution: *Suppose $V^* \equiv \max_i V_i > 0$, $\min_i V_i < V^*$, and $\frac{\Delta}{\kappa} < \ln n$. Then there exists a unique positive β_{ss} that satisfies (21). If we calculate $\mathcal{D}_{ss}$ and ρ_{ss} from (11) and (16) and find that $\rho_{ss} < \Delta^{-1}$, then β_{ss}, $\mathcal{D}_{ss}$, and ρ_{ss} solve (2), achieving the value $W_{ss} = \frac{1-\delta\Delta}{\beta_{ss}\delta\kappa}$.*

(c.) Immediate solution: *Suppose the conditions of part (b.) are satisfied, but the implied arrival probability in the first time step exceeds one: $\rho_{ss}\Delta \geq 1$. Then (2) is solved by the precision β^{Δ}_{ss} that satisfies (17), with entropy $\mathcal{D}^{\Delta}_{ss} = \Delta/\kappa$ and arrival rate $\rho^{\Delta}_{ss} = \Delta^{-1}$. Alternatively, suppose $\min_i V_i = \max_i V_i > 0$ or $\frac{\Delta}{\kappa} \geq \ln n$. Then (2) is solved with infinite precision and arrival rate $\rho^{\Delta}_{ss} = \Delta^{-1}$.*

[10] For computer simulations, numerical overflow is also a possibility, as discussed in the footnote to Prop. 1.

Figure 3: Optimal choice of rationality β when time discounting is the only cost of choice.

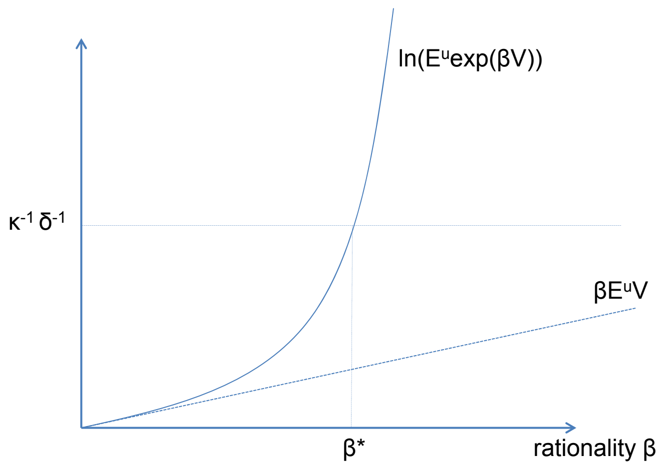

Note: The figure illustrates the interior solution stated in Prop. 2(b), as $\Delta \to 0$. Curve $g(\beta, \vec{V}) = \ln(E^u \exp(\beta V))$ is tangent to $\beta E^u V$ at $\beta = 0$, and has limiting slope $V^* \equiv \max_i V_i$ (assumed positive in graph) as $\beta \to \infty$.

If instead $V^* \equiv \max_i V_i < 0$, then $g(\beta, V)$ slopes downward, and it is optimal to never choose; see Prop. 2(a). If instead $\min_i V_i = \max_i V_i \equiv V^* > 0$, then $g(\beta, \vec{V})$ is a straight line from the origin, $g(\beta, \vec{V}) = \beta V^*$. Then the corner $\beta_{ss} = 0$ is preferred; see Prop. 2(c).

Again, the most relevant solution is the interior solution, which is unique when it exists. Postponement applies only in the trivial case when no available option has positive value. For a discrete-time numerical solution, one must allow for the possibility that the decision arrives in exactly one time step Δ. But in the continuous-time limit, the immediate solution arises only in the trivial case where all options give exactly the same value, so that $\min_i V_i = \max_i V_i$.

To summarize the comparative statics implications, first, Fig. 3 shows that an increase in the information cost parameter κ decreases β_{ss}. The expected value of the solved problem, $E^\pi V$, therefore decreases too, as does the information content of the decision, $K(\beta_{ss})$. The figure also shows that the proportional change in β_{ss} is smaller than the proportional change in κ. Hence an increase in κ increases the product $\kappa\beta_{ss}$, and so the problem value $W_{ss} = \frac{1-\delta\Delta}{\beta_{ss}\delta\kappa}$ falls.

Next, although the information content $K(\beta_{ss})$ increases with β_{ss}, the time cost of this information, $\kappa K(\beta_{ss})$, is subject to two offsetting effects. While an increase in κ decreases β_{ss} and thus $K(\beta_{ss})$, the product $\kappa K(\beta_{ss})$ is also increased by the direct impact of κ. Therefore the effect of κ on $\kappa K(\beta_{ss})$ appears to be ambiguous. Hence, the effect of κ on the arrival rate, $\rho_{ss} = \frac{1}{\kappa \mathcal{D}_{ss}}$, is also ambiguous.

By analogous reasoning, an increase in the discount rate δ decreases β_{ss}, $E^\pi V$, and W_{ss}. An increase in δ causes both $K(\beta_{ss})$ and $\kappa K(\beta_{ss})$ to decrease; hence ρ_{ss} rises with δ.

Convexity of the function g implies that a mean-preserving spread of the distribution of choice values $\vec{V}$ decreases the chosen β_{ss}. In Figure 3, a mean-preserving spread of $\vec{V}$ leaves the line $\beta E^u V$ unchanged, and shifts up the curve $g(\beta, \vec{V}) \equiv \ln(E^u \exp(\beta V))$, so β_{ss} must decrease. As discussed earlier, this model implicitly assumes that greater diversity in the values of the choice options $\vec{V}$ means they are objectively easier to distinguish. Thus, probabilities can be allocated more easily without increasing β.

2.2 One costly decision, with an alternative use of time

We next consider some extensions of the choice model from Section 2.1.[11] First, we allow for alternative uses of time, so that time devoted to decisions has additional costs, beyond pure discounting. We find that this slightly simplifies the classification of solutions, but there are still three cases: postponement, interior solution, or immediate solution.

Thus, let h be the fraction of time devoted to "work" or to any activity other than choice. Let the output from work during a time step Δ be $f(h)\Delta$, where $f' > 0$, $f'' \leq 0$, $\lim_{h\to 0} f'(h) = \infty$, and $\lim_{h\to\infty} f'(h) = 0$.[12] As before, let ρ represent the continuous-time arrival rate of the solution, so the expected time to completion is ρ^{-1}. But since the fraction of time devoted to calculation is only $1-h$, the expected time spent analyzing the problem before the solution arrives is only $\frac{1-h}{\rho}$. As before, we assume that the precision of the choice (measured by relative entropy) is limited by the expected time spent solving the problem. Thus the value W_t of this generalized problem is given by:

$$\begin{aligned}
W_t \quad &= \quad \max_{\rho,\, h,\, \{\pi_i\}_{i=1}^n} \quad f(h)\Delta + (1-\delta\Delta)\left[\rho\Delta\sum_{i=1}^n \pi_i V_i + (1-\rho\Delta)W_{t+\Delta}\right] \qquad (22) \\
\text{subject to:} \quad &\rho\kappa\left(\sum_{i=1}^n \pi_i \ln \pi_i + \ln n\right) \;\leq\; 1-h, \\
\text{and} \quad &\sum_{i=1}^n \pi_i = 1.
\end{aligned}$$

As before, let λ_ρ and λ_π represent the multipliers on the first and second constraints. At an interior solution with $\rho > 0$ and $0 < h < 1$, optimal tradeoffs between work, the arrival rate, and the probability allocation are summarized by

$$\lambda_\rho \;=\; f'(h)\Delta \;=\; \frac{(1-\delta\Delta)\rho\Delta V_i - \lambda_\pi}{\rho\kappa(1+\ln\pi_i)} \;=\; \frac{(1-\delta\Delta)\Delta(E^\pi V - W_{t+\Delta})}{\kappa\mathcal{D}(\vec{\pi}||\vec{u})}. \qquad (23)$$

Again, the first-order condition for π_i implies that probabilities take logit form:

$$\pi_i \;=\; \frac{\exp(\beta V_i)}{\sum_j \exp(\beta V_j)}, \qquad (24)$$

where now

$$\beta \;=\; \frac{1-\delta\Delta}{\kappa f'(h)}. \qquad (25)$$

[11] Sections 2.2-2.3, which discuss intermediate specifications, may be skipped without loss of continuity if the reader wishes to go directly to the main model, discussed in Sec. 2.4.

[12] We write the problem treating h as the fraction of a given time step devoted to work. One can instead write a model in which h represents the fraction of time steps in that are devoted entirely to work, with the others devoted entirely to calculation; this specification is also tractable and yields very similar results.

Total information use can then be calculated from the logit formula as

$$\mathcal{D} = \sum_i \pi_i \ln \pi_i + \ln n = \beta E^\pi V - \ln(E^u \exp(\beta V)). \tag{26}$$

But on the other hand, the first-order condition for the arrival rate ρ can be written as

$$\mathcal{D} \;=\; \beta(E^\pi V - W_{t+\Delta}). \tag{27}$$

Thus, combining the implications of optimal probabilities π_i and the optimal hazard ρ, we obtain the same equation for optimal precision that we found previously:

$$\beta W_{t+\Delta} \;=\; \ln(E^u \exp(\beta V)). \tag{28}$$

After calculating precision, we can back out the other variables from the remaining equilibrium equations; in particular, the arrival rate ρ must satisfy

$$1 - h \;=\; \rho\kappa\mathcal{D}. \tag{29}$$

Equation (28) is unchanged from Section 2.1. The only difference is that work is bounded between zero and one, so the range of precision that might be chosen in equilibrium is bounded bounded between $\frac{1-\delta\Delta}{\kappa f'(0)} = 0$ and $\beta_1^* \equiv \frac{1-\delta\Delta}{\kappa f'(1)}$. This slightly simplifies the classification of equilibrium types, since infinite precision is ruled out. Proposition 1 generalizes as follows.

Proposition 3 *Consider problem (22) when $W_{t+\Delta} \equiv W^+$ is a given parameter. Define $\beta_1^* \equiv \frac{1-\delta\Delta}{\kappa f'(1)}$, and $m(\beta) \equiv E^u \exp(\beta(V - W^+))$.*

(a.) Postponement: *Suppose $m(\beta_1^*) \leq 1$. Then it is optimal to postpone the solution, and instead devote the current time step to work only ($h^* = 1$), achieving the value $W_t^* = f'(1)\Delta + (1 - \delta\Delta)W^+$.*

(b.) Interior solution: *Suppose $V^* \equiv \max_i V_i > W^+ > E^u V$. Then there exists a unique positive β^* that satisfies (28). If $m(\beta_1^*) > 1$, then $\beta^* < \beta_1^*$. In this case, let $\mathcal{D}^*$, h^*, and ρ^* be given by (26), (25), and (29). If $\rho^* < \Delta^{-1}$, then β^*, $\mathcal{D}^*$, h^*, and ρ^* solve (22).*

(c.) Immediate solution: *Suppose $E^u V \geq W^+$, or that the conditions of part (b) are satisfied, but imply that the adjustment probability exceeds one in the first time step: $\rho^*\Delta \geq 1$. Then it is optimal to set the arrival rate $\rho_\Delta^* = \Delta^{-1}$. Precision and working time are given by the unique solution $(\beta_\Delta^*, \mathcal{D}^*, h_\Delta^*)$ of (26), (25), and (29) conditional on ρ_Δ^*.*

The postponement solution applies when $\max_i V_i \leq W^+$, because then there is no reason to devote valuable time to making the decision. In this case $m(\beta)$ slopes downward, and $m(\beta) < 1$ for all $\beta > 0$. But postponement is also optimal when $m(\beta) = 1$ is solved by a large positive number $\beta^* > \beta_1^*$. Such a solution means that the decision is only worthwhile if a high level of precision is achieved, and if $\beta^* > \beta_1^*$ this implies that the marginal time cost of achieving this precision is so high that the DM is better off devoting all time to work.

The interior solution $\beta^* > 0$ is characterized by $m(\beta^*) = 1$. The assumption $E^u V < W^+ < V^*$ implies that $m(\beta)$ initially slopes down from $m(0) = 1$, but is convex, and eventually increases without bound. Therefore the assumption $E^u V < W^+ < V^*$ implies that a unique positive β^* solves $m(\beta^*) = 1$. The only question, then, is whether that β^* is dominated by one of the corner solutions (postponement, or immediate solution). The former can be verified by checking whether $m(\beta_1^*) < 1$, which implies $\beta_1^* < \beta^*$, so that postponement is optimal, or $m(\beta_1^*) > 1$, which implies $\beta^* < \beta_1^*$, so that the interior solution is preferred to postponement. The latter can be verified by checking whether the arrival rate ρ^* associated with the interior solution β^* is less than Δ^{-1}; if not, an immediate one-step solution is optimal. There is then a unique solution for precision and hours worked, because (25) defines an upward-sloping relation between h and β, while (29) defines a downward-sloping relation between h and β, conditional on any ρ_Δ^*.

Note that in contrast to Props. 1-2, here it is irrelevant whether or not $\frac{\Delta}{\kappa}$ exceeds $\ln n$. When $\frac{\Delta}{\kappa} > \ln n$, a perfectly precise decision ($\beta^* = \infty$) can be achieved in a single time step Δ. This is optimal if time is only valuable as an input to the decision problem. But when time has a nonzero marginal value in other uses, perfect precision cannot be optimal. Instead, the optimal tradeoff between decision-making and other uses of time is characterized by (25), which shows that $\beta^* < \infty$ as long as $f'(h) > 0$ and $\kappa > 0$.

Next, we can also consider a "steady state" solution for $W_t = W_{t+\Delta} \equiv W_{ss}$ in an infinite-horizon, time-invariant version of (22). The steady state values β_{ss}, h_{ss}, and W_{ss} are characterized by (25), (28), and

$$W_{ss} \;=\; \frac{f(h_{ss}) + (1-h_{ss})f'(h_{ss})}{\delta}. \tag{30}$$

The steady-state solution can be described as follows.

Proposition 4 *Consider a time-invariant, infinite-horizon version of problem (22). Call the value $W_{ss} \equiv W_t$ for any t, and define $\beta(h) \equiv \frac{1-\delta\Delta}{\kappa f'(h)}$.*

(a.) Postponement: *Suppose $\frac{\ln E^u \exp(\beta(1)V)}{\beta(1)} \leq \frac{f(1)}{\delta}$. Then it is optimal to postpone solving, and instead devote all time to work ($h_{ss} = 1$), achieving the value $W_{ss} = \frac{f(1)}{\delta}$.*

(b.) Interior solution: *Suppose $\frac{\ln E^u \exp(\beta(1)V)}{\beta(1)} > \frac{f(1)}{\delta}$. Then there exists a unique $h_{ss} \in (0,1)$ that satisfies $\frac{\ln E^u \exp(\beta_{ss}V)}{\beta_{ss}} = \frac{f(h_{ss})+(1-h_{ss})f'(h_{ss})}{\delta}$, where $\beta_{ss} \equiv \beta(h_{ss})$. If the associated arrival rate $\rho_{ss} \equiv \frac{1-h_{ss}}{\kappa K(\beta_{ss})}$ is less than Δ^{-1}, then h_{ss}, β_{ss}, and ρ_{ss} solve (22).*

(c.) Immediate solution: *Suppose $\frac{\ln E^u \exp(\beta(1)V)}{\beta(1)} > \frac{f(1)}{\delta}$, but the rate ρ_{ss} defined in (b) exceeds Δ^{-1}. Then it is optimal to set the arrival rate $\rho_{ss}^\Delta \equiv \Delta^{-1}$. Optimal time worked is the unique h_{ss}^Δ that satisfies $\Delta^{-1} = \frac{1-h_{ss}^\Delta}{\kappa K(\beta_{ss}^\Delta)}$, where $\beta_{ss}^\Delta \equiv \beta(h_{ss}^\Delta)$ is the optimal precision.*

Proof. We have assumed $\lim_{h\to 0} f'(h) = \infty$. Equation (30) defines W as a decreasing function of h, with an infinite limit at $h = 0$. Equation (25) defines a relation $\beta(h)$ which is increasing in h, and in turn we can calculate W from β from the equation

$$W \;=\; \frac{\ln(E^u \exp(\beta V))}{\beta}. \tag{31}$$

Equation (31) defines W as a function of β; the limiting values of the curve are $E^u V$ at $\beta = 0$, and $V^* = \max_i V_i$ as $\beta \to \infty$. Differentiating, it can be shown that the slope of the curve (31) equals $\frac{\beta'(h)K(\beta)}{(\beta(h))^2}$. Since $\beta'(h) > 0$ and $K(\beta)$ represents a Kullback-Leibler divergence, which is nonnegative by construction, (31) represents an increasing function of h. Thus (30) and (31) can cross at most once. If they cross at a value below $h = 1$, then (22) has an interior solution; otherwise, (22) has a corner solution, as described in the proposition. **QED**.

2.3 Choosing a stopping time

As Propositions 1-4 make clear, given the option to choose among several alternatives i with values V_i, it may be optimal to ignore this opportunity altogether, or to choose as quickly as possible, or to choose with an intermediate degree of precision. But this in itself represents a choice across several alternatives— concretely, it represents a stopping time problem, determining *when* the solution of the choice across alternatives V_i will arrive. This stopping time problem was solved perfectly optimally in the preceding sections. But in the present context it is more natural to assume that timing decisions are subject to error as well. Therefore in this section we study the costly, error-prone choice of a stopping time.[13] That is, we assume that sooner or later the DM may wish to exercise an option with (possibly time-varying) value X. But the DM must also decide *when* to exercise this option, and this timing decision is costly. As in Sec. 2.1, we also allow for an alternative use of time, besides decision-making.

We model the choice of *when* to exercise an option as the choice of a *hazard rate*, ρ, at which the decision process stops and the option is exercised. In analogy to our previous setup, we assume that choosing the hazard rate is costly, and that the cost increases with precision. Previously, we interpreted "imprecise" choice as a uniform distribution over a set of alternatives; here, we interpret "imprecise" choice of the hazard rate as a uniform hazard with an exogenous rate $\bar{\rho}$, which is a free parameter in our framework. Cost increases with deviations from the actual hazard from the uniform hazard at rate $\bar{\rho}$, as measured by Kullback-Leibler divergence.[14]

Let h be time spent "working", and μ be time spent "monitoring" the situation, to decide whether or not to stop. Choice of the stopping hazard is governed by the following Bellman equation.

$$\begin{aligned} W_t \quad &= \quad \max_{h,\mu,\rho} \; f(h)\Delta + (1-\delta\Delta)\left[\rho\Delta X_{t+\Delta} + (1-\rho\Delta)W_{t+\Delta}\right] \qquad (32) \\ \text{subject to:} \quad &\kappa\mathcal{D}\left(\left(\begin{array}{c}\rho\Delta \\ 1-\rho\Delta\end{array}\right)\middle|\middle|\left(\begin{array}{c}\bar{\rho}\Delta \\ 1-\bar{\rho}\Delta\end{array}\right)\right) \;\le\; \mu\Delta \quad \text{and} \quad h+\mu\le 1. \end{aligned}$$

The Kullback-Leibler divergence that appears on the left-hand side of the constraint is given by

$$\mathcal{D}\left(\left(\begin{array}{c}\rho\Delta \\ 1-\rho\Delta\end{array}\right)\middle|\middle|\left(\begin{array}{c}\bar{\rho}\Delta \\ 1-\bar{\rho}\Delta\end{array}\right)\right) \;=\; \rho\Delta\ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho\Delta)\ln\left(\frac{1-\rho\Delta}{1-\bar{\rho\Delta}}\right). \qquad (33)$$

In the continuous-time limit $\Delta \to 0$, the precision constraint simplifies to

$$\kappa\rho\ln\left(\frac{\rho}{\bar{\rho}}\right) \;\le\; \mu. \qquad (34)$$

But to avoid continuous-time technicalities, we will solve the problem assuming a small discrete time step Δ. We will then see that the discrete-time model has a well-defined and well-behaved limit as we go to continuous time.

[13] The error-prone stopping time decision modeled in this subsection is analogous to the "Woodford-control" specification studied in Costain and Nakov (2014).

[14] At first glance, it might seem more natural to treat the decision to adjust or not to adjust in a single time step as a binary decision, applying the model from the previous subsections with uniform benchmark $\vec{u} \equiv \{0.5, 0.5\}$ on the two options. But such a model would not be well behaved in the limit as $\Delta \to 0$. The model would imply a 50% probability of adjustment in one time step, conditional on indifference between adjustment and nonadjustment, *regardless of* Δ. In other words, as $\Delta \to 0$ such a model would imply perfectly rational timing, regardless of κ. The introduction of the benchmark adjustment hazard $\bar{\rho}$ is what gives our model a well-behaved continuous-time limit where timing errors still occur.

If we denote the multipliers on the constraints by λ_μ and λ_h, the first-order conditions are

$$\begin{aligned}
f'(h) - \lambda_h &= 0 && (35)\\
\lambda_\mu \Delta - \lambda_h &= 0 && (36)\\
(1-\delta\Delta)(X_{t+\Delta} - W_{t+\Delta}) - \lambda_\mu \kappa \left(\ln\left(\frac{\rho}{\bar{\rho}}\right) + 1 - \ln\left(\frac{1-\rho\Delta}{1-\bar{\rho}\Delta}\right) - 1 \right) &= 0 && (37)\\
\mu\Delta - \kappa\mathcal{D} &= 0 && (38)\\
1 - \mu - h &= 0. && (39)
\end{aligned}$$

Eliminating the multipliers and rearranging, we obtain

$$\ln\left(\frac{\rho\Delta}{1-\rho\Delta}\right) = \ln\left(\frac{\bar{\rho}\Delta}{1-\bar{\rho}\Delta}\right) + \beta(X_{t+\Delta} - W_{t+\Delta}) \tag{40}$$

where

$$\beta = \frac{1-\delta\Delta}{\kappa f'(h)}. \tag{41}$$

Solving for the endogenous hazard ρ, we obtain

$$\begin{aligned}
\rho &= \frac{\bar{\rho}\exp(\beta X_{t+\Delta})}{(1-\bar{\rho}\Delta)\exp(\beta W_{t+\Delta}) + \bar{\rho}\Delta\exp(\beta X_{t+\Delta})} && (42)\\
&= \frac{\bar{\rho}\exp(\beta D_{t+\Delta})}{1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta D_{t+\Delta})} && (43)\\
&= \frac{\bar{\rho}}{\bar{\rho}\Delta + (1-\bar{\rho}\Delta)\exp(-\beta D_{t+\Delta})}, && (44)
\end{aligned}$$

where $D_{t+\Delta} = X_{t+\Delta} - W_{t+\Delta}$. This equation expresses the stopping probability $\rho\Delta$ as a weighted binomial logit that places weight $\bar{\rho}\Delta$ on stopping, and weight $1-\bar{\rho}\Delta$ on continuation. Note that as the precision β approaches zero, the process stops with a constant hazard $\bar{\rho}$ regardless of the value of stopping, while as $\beta \to \infty$, the process stops if and only if $X_{t+\Delta} > W_{t+\Delta}$. The parameter $\bar{\rho}$ controls the *speed* of stopping; it represents the stopping hazard when the DM is indifferent between stopping and continuing. Notice also that the hazard rate simplifies as we go to continuous time, converging to

$$\rho = \bar{\rho}\exp(\beta D_{t+\Delta}) \tag{45}$$

in the limit as $\Delta \to 0$. Note also that under this logit probability, (33) simplifies to

$$\begin{aligned}
\mathcal{D} &= \rho\Delta\beta D_{t+\Delta} - \ln(1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta D_{t+\Delta})) && (46)\\
&= \frac{\bar{\rho}\Delta\beta D_{t+\Delta}}{\bar{\rho}\Delta + (1-\bar{\rho}\Delta)\exp(-\beta(D_{t+\Delta}))} - \ln(1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta D_{t+\Delta})) \equiv K_\rho(\beta). && (47)
\end{aligned}$$

The function $K_\rho(\beta)$ represents a cost function for precision. It is straightforward to show by differentiation that $K'_\rho(\beta) > 0$.

The expected gains g from the option to stop, per period, are

$$g \equiv \rho(1-\delta\Delta)D_{t+\Delta} = \frac{\bar{\rho}(1-\delta\Delta)D_{t+\Delta}}{\bar{\rho}\Delta + (1-\bar{\rho}\Delta)\exp(-\beta(D_{t+\Delta}))}. \tag{48}$$

Alternatively, we can calculate the *net* expected gains G, as follows:

$$G \equiv \rho(1-\delta\Delta)D_{t+\Delta} - \kappa f'(h)\mathcal{D} = \kappa f'(h) \ln\left(1 - \bar{\rho}\Delta + \bar{\rho}\Delta \exp\left(\frac{(1-\delta\Delta)D_{t+\Delta}}{\kappa f'(h)}\right)\right). \quad (49)$$

Here ρD represents the gross gains from stopping (in units of goods); $\kappa f'(h)\mathcal{D}$ represents the costs of monitoring the situation, where $\kappa\mathcal{D}$ is in units of time, and $f'(h)$ converts time into goods units. Note that (49) corresponds to equation (22) of Costain and Nakov (2014).

The models considered in Secs. 2.1-2.2 had a variety of possible corner solutions, in which the decision process ended with probability zero or one in a single time step. In contrast, the present model penalizes stopping with probability zero or one. This simplifies the structure of this decision problem, ensuring an interior solution for precision at any step of the backwards induction process defined by (32). To see this, write the time constraint as

$$1 - h = \mu = \frac{\kappa}{\Delta} K_\rho(\beta), \quad (50)$$

where $K_\rho(\beta)$ is given by (47).

Proposition 5 *Define $D_{t+\Delta} \equiv X_{t+\Delta} - W_{t+\Delta}$, where $X_{t+\Delta}$ and $W_{t+\Delta}$ are given parameters. Suppose $\lim_{h\to 0} f'(h) = \infty$. Then there is a unique pair (h^*, β^*), with $\beta^* > 0$ and $h^* \in (0,1)$, that satisfy (41) and (47). The pair (h^*, β^*) solves the problem (32), given $X_{t+\Delta}$ and $W_{t+\Delta}$.*

Proof. Note that for $h \in [0,1]$, the first-order condition (41) defines β as an increasing function of h. Since $\lim_{h\to 0} f'(h) = \infty$, the function starts at zero, and it increases continuously to $\beta_h^* \equiv \frac{1-\delta\Delta}{\kappa f'(1)}$. On the other hand, since $K_\rho(\beta)$ is an increasing function, (50) is a decreasing function that reaches $\beta = 0$ at $h = 1$ (when no time is devoted to the decision). Therefore there is a unique interior solution for β and h. **QED.**[15]

This proposition is the analog of Prop. 1, showing that there is a unique solution to a single step of the iteration of the Bellman equation. It is also interesting to solve for the steady state of the Bellman equation. That is, fixing a time-independent value $X_{t+\Delta} \equiv X_{ss}$ of the stopped process, we wish to solve for the limit $W_{ss} \equiv W_t = W_{t+\Delta}$ as we iterate the backwards induction problem. From the Bellman equation, the value W_{ss} of the unstopped process must satisfy

$$W_{ss}(\delta + (1-\delta\Delta)\rho_{ss}) = f(h_{ss}) + (1-\delta\Delta)\rho_{ss}X_{ss}. \quad (51)$$

The steady state arrival rate ρ_{ss} must solve (44), where

$$D_{ss} = X_{ss} - W_{ss} = \frac{(\delta + (1-\delta\Delta)\rho_{ss})X_{ss} - f(h_{ss}) - (1-\delta\Delta)\rho_{ss}X_{ss}}{\delta + \rho_{ss} - \delta\rho_{ss}\Delta} = \frac{\delta X_{ss} - f(h_{ss})}{\delta + (1-\delta\Delta)\rho_{ss}}. \quad (52)$$

Substituting in for the arrival rate ρ_{ss}, we obtain

$$D_{ss} = \frac{\delta X_{ss} - f(h_{ss})}{\delta + (1-\delta\Delta)\frac{\bar{\rho}}{\bar{\rho}\Delta + (1-\bar{\rho}\Delta)\exp(-\beta_{ss}D_{ss})}}. \quad (53)$$

[15]If $X_{t+\Delta} = W_{t+\Delta}$ exactly, then $K_\rho(\beta) = 0$ for all $\beta \geq 0$, and the solution is $h^* = 1$ and $\rho^* = \bar{\rho}$, with β^* undefined. But as long as there is any difference between X and W, then the solution is strictly interior, with $\beta^* > 0$, $h^* < 1$, and $\rho^* \in (0, \Delta^{-1})$.

On the other hand, substituting the cost function $K_\rho(\beta_{ss})$ into (50), we have

$$1 - h_{ss} \;=\; \frac{\kappa\bar{\rho}\beta_{ss}D_{ss}}{\bar{\rho}\Delta + (1-\bar{\rho}\Delta)\exp(-\beta D_{ss})} - \frac{\kappa}{\Delta}\ln(1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta_{ss}D_{ss})). \tag{54}$$

Considering the fact that β is a function of h, equations (53) and (54) are two equations to determine h_{ss} and D_{ss}. To simplify, define $D_\beta \equiv \beta D_{t+\Delta}$, and consider the limit as $\Delta \to 0$. Then (53)-(54) become

$$\kappa f'(h)D_\beta \;=\; \frac{\delta X^+ - f(h)}{\delta + \bar{\rho}\exp(D_\beta)}, \tag{55}$$

$$1 - h \;=\; \bar{\rho}\kappa D_\beta(\exp(D_\beta) - 1). \tag{56}$$

Equation (55) is a relation between the value of the problem, W^+ (and hence the gain $D_{t+\Delta} = X^+ - W^+$), and the arrival rate of the solution, ρ. Equation (56) relates the time devoted to solving the problem, $1-h$, to the accuracy of the solution, as summarized by $\beta D_{t+\Delta}$.

With some additional algebra, (55) can be replaced by

$$D_\beta \;=\; \frac{\delta X^+ - f(h) - (1-h)f'(h)}{(\bar{\rho}+\delta)\kappa f'(h)}. \tag{57}$$

We can then use (56)-(57) to characterize the steady-state solution to (32).

Proposition 6 *Consider a steady-state solution to problem (32), taking $X_{t+\Delta} \equiv X^+$ as given, and defining $W_t = W_{t+\Delta} \equiv W$. Suppose $\lim_{h\to 0} f'(h) = \infty$.*

a. *A solution to (32) exists, and satisfies equations (56)-(57). This solution has $h^* > 0$, $\beta^* > 0$, and $\rho^*\Delta \in (0,1)$.*

b. *Suppose $X^+ \leq 0$. Then there is only one solution to (56)-(57). At this solution, $h^* > 0$, $\beta^* > 0$, $\rho^*\Delta \in (0,1)$, and $W > X^+$. This solution varies smoothly with small parameter changes.*

c. *Suppose $X^+ \geq \frac{f(1)}{\delta}$. Then there is only one solution to (56)-(57). At this solution, $h^* > 0$, $\beta^* > 0$, $\rho^*\Delta \in (0,1)$, and $W < X^+$. This solution varies smoothly with small parameter changes.*

Thus, a steady-state solution to the stopping time problem exists, and can be found by solving (56)-(57). There is only one solution to (56)-(57) when X^+ is sufficiently large. In this case, the decision-maker is impatient to stop, and the value W of the as-yet-unstopped problem is less than X^+. Likewise, there is only one solution to (56)-(57) if X^+ is negative. In this case, the decision-maker wishes to avoid stopping, and the value W of the as-yet-unstopped problem exceeds X^+. For intermediate values of X^+, in $\left(0, \frac{f(1)}{\delta}\right)$, the solution is also unique, but the proof of this fact will be postponed to the more general model of the next subsection.

2.4 Choosing when to solve a problem

Next, we combine the two decisions we have analyzed thus far: a DM chooses a stopping time that determines when he will solve a problem. As in Section 2.1, we assume that the DM chooses an arrival rate ρ for the solution, subject to a constraint on the precision of the decision relative to the expected time spent solving the problem. But now, rather than assuming that ρ is chosen optimally, we assume that the precision of the choice of the arrival rate ρ is constrained too, as in Sec. 2.3. Errors in the setting of the arrival rate amount to errors in the timing of the solution of the problem.[16]

In contrast with the previous models, we consider three possible uses of time. The fraction of time dedicated to solving the decision problem is called τ; the fraction dedicated to some alternative activity, such as work or leisure, is called h. Finally, let μ be the fraction of time dedicated to "monitoring" whether the current moment is a good time to solve the problem.

To understand μ, note that in a given time step of length Δ, the decision maker can choose an arrival probability $\rho\Delta \in [0,1]$, which amounts to choosing a hazard rate $\rho \in [0, \Delta^{-1}]$. In contrast with Sections 2.1-2.2, we allow for possible errors in the choice of the probability $\rho\Delta$. We measure the cost of this timing decision in terms of the relative entropy of the probabilities $(\rho\Delta, 1-\rho\Delta)$. Since the decision maker is effectively choosing a *hazard rate*, we measure entropy relative to probabilities $(\bar{\rho}\Delta, 1-\bar{\rho}\Delta)$, where $\bar{\rho}$ is some underlying hazard rate, which is a free parameter of the model. Assuming that the fraction of the time step dedicated to this decision is μ, the total time dedicated to the decision is $\mu\Delta$. Therefore the constraint on the precision of the timing decision takes the form

$$\kappa_\rho \mathcal{D}\left((\rho\Delta, 1-\rho\Delta)\,||\,(\bar{\rho}\Delta, 1-\bar{\rho}\Delta)\right) \;\leq\; \mu\Delta. \tag{58}$$

Note that in the continuous-time limit ($\Delta \to 0$), the constraint on the precision of timing reduces to

$$\kappa_\rho \rho \ln\left(\frac{\rho}{\bar{\rho}}\right) \;\leq\; \mu. \tag{59}$$

For brevity, the constraint is shown in this continuous-time form in Figure 4, which illustrates both margins of this model of costly decision in extensive form. However, to avoid continuous-time technicalities, we solve the problem in its discrete-time formulation below.

Assuming a discrete time step Δ, the value W_t will be governed the following Bellman equation:

$$\begin{aligned}
W_t \quad = & \max_{\rho, h, \tau, \mu, \{\pi_i\}_{i=1}^n} \quad f(h)\Delta + (1-\delta\Delta)W_{t+\Delta} + \rho\Delta(1-\delta\Delta)\left(\sum_{i=1}^n \pi_i V_i - W_{t+\Delta}\right) \\
\text{s.t.:} \quad & \rho\kappa_\pi \left(\sum_{i=1}^n \pi_i \ln \pi_i + \ln n\right) \;\leq\; \tau, \\
\text{and} \quad & \kappa_\rho\left(\rho\Delta \ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho\Delta)\ln\left(\frac{1-\rho\Delta}{1-\bar{\rho}\Delta}\right)\right) \;\leq\; \mu\Delta, \\
\text{and} \quad & \sum_{i=1}^n \pi_i = 1, \quad \text{and} \quad h + \mu + \tau \leq 1.
\end{aligned} \tag{60}$$

In the first constraint, the decision costs $\kappa_\pi \mathcal{D}(\vec{\pi}||\vec{u})$ are constrained by τ/ρ, which is the expected time dedicated to the decision prior to its arrival (given arrival rate ρ, and fraction τ of time dedicated to the decision).

[16] The model developed in this subsection is analogous to the "Nested-control" specification analyzed by Costain and Nakov (2014).

Figure 4: Choosing when to solve a problem.

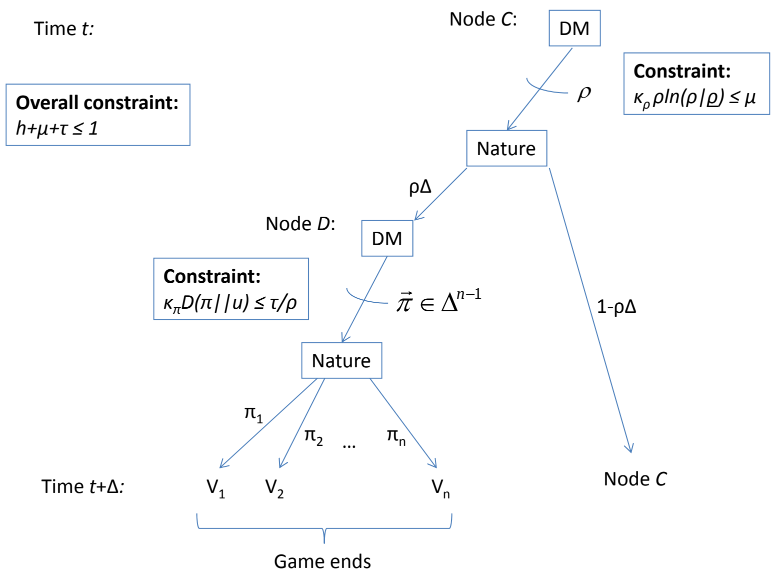

Note: Decision maker (DM) chooses arrival rate ρ of decision. Nature lets decision arrive with probability $\rho\Delta$. A slower decision implies higher precision in the allocation of probabilities π_i across alternatives i with values V_i, satisfying the constraint $\kappa_\pi \rho \mathcal{D}(\vec{\pi}||\vec{u}) \leq \tau$, where τ is the fraction of time devoted to the choice between alternatives V_i. Choosing the arrival rate ρ is also costly, and must satisfy the constraint $\kappa_\rho \rho \ln(\rho/\bar{\rho}) \leq \mu$, where μ is the fraction of time devoted to the timing decision.

We will write the multipliers on the constraints as λ_τ, λ_μ, λ_π, and λ_h, respectively. The decision maker equalizes the marginal value of the three uses of time, which implies the following first-order conditions:

$$\lambda_h = f'(h)\Delta = \lambda_\tau = \lambda_\mu \Delta. \tag{61}$$

The remaining first-order conditions can be written as

$$(1-\delta\Delta)\rho\Delta V_i - \lambda_\tau \rho \kappa_\pi (1 + \ln \pi_i) - \lambda_\pi \;=\; 0 \tag{62}$$

$$(1-\delta\Delta)\Delta\left(\sum_{i=1}^{n} \pi_i V_i - W_{t+\Delta}\right) - \lambda_\tau \kappa_\pi \mathcal{D}_\pi - \lambda_\mu \Delta \kappa_\rho \left[\ln\left(\frac{\rho}{\bar{\rho}}\right) + 1 - \ln\left(\frac{1-\rho\Delta}{1-\bar{\rho}\Delta}\right) - 1\right] \;=\; 0 \tag{63}$$

$$\tau - \rho\kappa_\pi \mathcal{D}_\pi \;=\; 0 \tag{64}$$

$$\mu\Delta - \kappa_\rho \mathcal{D}_\rho \;=\; 0 \tag{65}$$

$$1 - h - \mu - \tau \;=\; 0 \tag{66}$$

$$1 - \sum_{i=1}^{n} \pi_i \;=\; 0 \tag{67}$$

where

$$\mathcal{D}_\pi \;\equiv\; \sum_{i=1}^{n} \pi_i \ln \pi_i + \ln n \tag{68}$$

$$\mathcal{D}_\rho \;\equiv\; \rho\Delta \ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho\Delta)\ln\left(\frac{1-\rho\Delta}{1-\bar{\rho}\Delta}\right). \tag{69}$$

As before, (62) and (67) show that the optimal probabilities π_i take the form of a logit, as in Costain and Nakov (2014), equation (14):

$$\pi_i \;=\; \frac{\exp(\beta_\pi V_i)}{\sum_j \exp(\beta_\pi V_j)}, \tag{70}$$

where

$$\beta_\pi \;=\; \frac{1-\delta\Delta}{\kappa_\pi f'(h)}. \tag{71}$$

We can follow our previous calculations to find an explicit formula for the information $\mathcal{D}_\pi$ used in choosing these probabilities:

$$\mathcal{D}_\pi \;=\; \beta_\pi E^\pi V - \ln(E^u \exp(\beta_\pi V)) \;=\; \beta_\pi \frac{E^u V \exp(\beta_\pi V)}{E^u \exp(\beta_\pi V)} - \ln(E^u \exp(\beta_\pi V)). \tag{72}$$

Condition (63) for the optimal arrival probability ρ can be rewritten as

$$\frac{\kappa_\rho f'(h)}{1-\delta\Delta}\ln\left(\frac{\rho/(1-\rho\Delta)}{\bar{\rho}(1-\bar{\rho}\Delta)}\right) \;=\; E^\pi V - W_{t+\Delta} - \frac{\kappa_\pi f'(h)}{1-\delta\Delta}\mathcal{D}_\pi\,, \tag{73}$$

which simplifies to

$$\frac{\rho\Delta}{1-\rho\Delta} \;=\; \frac{\bar{\rho}\Delta}{1-\bar{\rho}\Delta}\exp\left(\beta_\rho D_{t+\Delta}\right), \tag{74}$$

where

$$\beta_\rho \;\equiv\; \frac{1-\delta\Delta}{\kappa_\rho f'(h)} \;=\; \frac{\kappa_\pi}{\kappa_\rho}\beta_\pi, \tag{75}$$

$$D_{t+\Delta} \;\equiv\; \tilde{V}_{t+\Delta} - W_{t+\Delta}, \tag{76}$$

$$\tilde{V}_{t+\Delta} \;\equiv\; E^\pi V - \frac{\mathcal{D}_\pi}{\beta_\pi}. \tag{77}$$

Equation (75) represents an optimal tradeoff between the allocation of precision (and time) to the decision about *when* to adjust, and to the decision about *which* option to choose, when adjusting. The quantity $\tilde{V}_{t+\Delta}$ represents the expected gains that accrue upon adjustment, net of adjustment costs; the factor β_π converts precision into time at rate κ_π, time into utility at rate $f'(h)$, and next period's utility into utility now at rate $1-\delta\Delta$, so that the payoff $E^\pi V$ is commensurate with the precision measure $\mathcal{D}_\pi$. Using (72), we can calculate

$$\tilde{V}_{t+\Delta} \;=\; \beta_\pi^{-1}\ln\left(E^u \exp\left(\beta_\pi V\right)\right). \tag{78}$$

This corresponds to the explicit formula for the value function given in Costain and Nakov (2014), section 2.2.1.

Rearranging the first-order condition (74), the solution arrives with probability

$$\rho\Delta = \frac{\bar{\rho}\Delta}{\bar{\rho}\Delta + (1-\bar{\rho}\Delta)\exp\left(-\beta_\rho D_{t+\Delta}\right)} \tag{79}$$

$$= \frac{\bar{\rho}\Delta \exp\left(\beta_\rho \tilde{V}_{t+\Delta}\right)}{\bar{\rho}\Delta \exp\left(\beta_\rho \tilde{V}_{t+\Delta}\right) + (1-\bar{\rho}\Delta)\exp\left(\beta_\rho W_{t+\Delta}\right)} \tag{80}$$

$$= \frac{\bar{\rho}\Delta\left(E^u \exp(\beta_\pi V)\right)^{\frac{\kappa_\pi}{\kappa_\rho}}}{\bar{\rho}\Delta\left(E^u \exp(\beta_\pi V)\right)^{\frac{\kappa_\pi}{\kappa_\rho}} + (1-\bar{\rho}\Delta)\left(\exp\left(\beta_\pi W_{t+\Delta}\right)\right)^{\frac{\kappa_\pi}{\kappa_\rho}}}. \tag{81}$$

This is a weighted binary logit that compares the values of adjustment and nonadjustment; it clearly implies $\rho\Delta \in [0,1]$. If we assume that the time unit is the same as the time step in the game (implying $\Delta \equiv 1$), then it corresponds to equation (20) of Costain and Nakov (2014). The formula also has a simple continuous-time limit; as $\Delta \to 0$, we have

$$\rho \to \bar{\rho}\exp\left(\beta_\rho D_{t+\Delta}\right) = \bar{\rho}\exp\left(\beta_\rho(\tilde{V}_{t+\Delta} - W_{t+\Delta})\right). \tag{82}$$

Note therefore that the continuous-time limit offers very simple parallel ways of writing the first-order conditions for the probability of adjustment and the probabilities of choosing each possible option:

$$\ln \pi_i = \ln(1/n) + \beta_\pi(V_i - \tilde{V}_{t+\Delta}), \tag{83}$$

$$\ln \rho = \ln \bar{\rho} + \beta_\rho(\tilde{V}_{t+\Delta} - W_{t+\Delta}). \tag{84}$$

Finally, consider the time devoted to choosing the adjustment timing. Given the formula for ρ, we now have $\ln(\rho/\bar{\rho}) = \beta_\rho D_{t+\Delta} - \ln(1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta_\rho D_{t+\Delta}))$, and $\ln((1-\rho\Delta)/(1-\bar{\rho}\Delta)) = -\ln(1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta_\rho D_{t+\Delta}))$. Therefore, (69) simplifies to

$$\mathcal{D}_\rho = \rho\Delta\beta_\rho D_{t+\Delta} - \ln(1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta_\rho D_{t+\Delta})) \tag{85}$$

$$= \frac{\bar{\rho}\Delta\beta_\rho D_{t+\Delta}}{\bar{\rho}\Delta + (1-\bar{\rho}\Delta)\exp\left(-\beta_\rho D_{t+\Delta}\right)} - \ln(1-\bar{\rho}\Delta + \bar{\rho}\Delta\exp(\beta_\rho D_{t+\Delta})). \tag{86}$$

Using L'Hopital's rule, this simplifies in the continuous-time limit to

$$\lim_{\Delta\to 0}\frac{\mathcal{D}_\rho}{\Delta} = \rho\beta_\rho D_{t+\Delta} + \bar{\rho}(1-\exp(\beta_\rho D_{t+\Delta})) = \rho\left(\beta_\rho D_{t+\Delta} - 1\right) + \bar{\rho}. \tag{87}$$

2.4.1 Backwards induction

As long as problem (60) has a well-defined maximum, it is easy to show that it defines a contraction. We begin by proving this fact, and then go on to show that the maximum exists under very weak assumptions.

Proposition 7 *Taking as given the continuation value $W_{t+\Delta}$, assume that the maximum in (60) exists. Then:*
(a.) the operator defined by problem (60) is a contraction, and
(b.) if the problem is defined over an infinite horizon and the options $i = 1..n$ have time-independent values V^i_{ss}, then (60) has a unique time-independent value $W_t = W_{t+\Delta} \equiv W_{ss}$.

Proof. Consider two possible continuation values $W^a_{t+\Delta} < W^b_{t+\Delta}$. Let $(h^a, \rho^a, \vec{\pi}^a)$ be the policy that solves (60) when $W_{t+\Delta} = W^a_{t+\Delta}$, with time t value W^a_t; and let $(h^b, \rho^b, \vec{\pi}^b)$ be the solution when $W_{t+\Delta} = W^b_{t+\Delta} \equiv W^a_{t+\Delta} + 1$, implying time t value W^b_t. Then we have

$$
\begin{aligned}
W^+ &\equiv f(h^b)\Delta + (1-\delta\Delta)W^a_{t+\Delta} + \rho^b\Delta(1-\delta\Delta)(E^{\pi^b}V - W^a_{t+\Delta}) + (1-\delta\Delta) && (88)\\
&\geq f(h^b)\Delta + (1-\delta\Delta)W^a_{t+\Delta} + \rho^b\Delta(1-\delta\Delta)(E^{\pi^b}V - W^a_{t+\Delta}) + (1-\rho^b\Delta)(1-\delta\Delta) && (89)\\
= W^b_t &= f(h^b)\Delta + (1-\delta\Delta)W^b_{t+\Delta} + \rho^b\Delta(1-\delta\Delta)(E^{\pi^b}V - W^b_{t+\Delta}) && (90)\\
&\geq f(h^a)\Delta + (1-\delta\Delta)W^b_{t+\Delta} + \rho^a\Delta(1-\delta\Delta)(E^{\pi^a}V - W^b_{t+\Delta}) && (91)\\
&= f(h^a)\Delta + (1-\delta\Delta)W^a_{t+\Delta} + \rho^a\Delta(1-\delta\Delta)(E^{\pi^a}V - W^a_{t+\Delta}) + (1-\rho^a\Delta)(1-\delta\Delta) && (92)\\
\geq W^a_t &= f(h^a)\Delta + (1-\delta\Delta)W^a_{t+\Delta} + \rho^a\Delta(1-\delta\Delta)(E^{\pi^a}V - W^a_{t+\Delta}) && (93)\\
\geq W^-_t &= f(h^b)\Delta + (1-\delta\Delta)W^a_{t+\Delta} + \rho^b\Delta(1-\delta\Delta)(E^{\pi^b}V - W^a_{t+\Delta}). && (94)
\end{aligned}
$$

Note that

$$W^b_t - W^a_t \;\leq\; (1-\delta\Delta) \;=\; W^+ - W^-,$$

whereas

$$W^b_{t+\Delta} - W^a_{t+\Delta} = 1.$$

Thus (60) defines a contraction of modulus $(1-\delta\Delta)$. Part (b) follows as an immediate corollary. **QED**.

Now, to see that (60) has a well-defined maximum, note that conditional on the values $\vec{V}_{t+\Delta}$ and $W_{t+\Delta}$ associated with the time $t+\Delta$ problem, the time t decision can be reduced to two equations. Specifically, the time constraint must bind, and the first-order condition for precision must be satisfied:

$$
\begin{aligned}
1-h &= \rho\kappa_\pi\mathcal{D}_\pi + \kappa_\rho\frac{\mathcal{D}_\rho}{\Delta}, && (95)\\
\kappa_\pi\beta_\pi f'(h) &= 1-\delta\Delta. && (96)
\end{aligned}
$$

These two equations (illustrated by Figs. 5-6 in the next subsection) effectively depend on only two variables, h and β_π, because all other time-t endogenous variables can easily be substituted out. $\mathcal{D}_\pi$ can be eliminated using equation (72), and ρ can be eliminated using (79) or (81). If we also calculate β_ρ and $\tilde{V}_{t+\Delta}$ using (75) and (78), we can then eliminate $\mathcal{D}_\rho$ using (85).

The first-order condition (96) clearly gives h as an increasing function of β_π. It is therefore helpful to analyze the slope of the time constraint (95), viewed as a relation between β_π and h. It will be helpful to use the identities $\beta_\pi D_{t+\Delta} = \ln(E^u\exp(\beta_\pi(V - W_{t+\Delta})))$ and $\exp\left(\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+\Delta}\right) = (E^u\exp(\beta_\pi(V - W_{t+\Delta})))^{\kappa_\pi/\kappa_\rho}$. Substituting, the time constraint can be written as follows:

$$
\begin{aligned}
1-h &= \rho(\beta_\pi)\kappa_\pi\left(\beta_\pi\frac{E^uVe^{\beta_\pi V}}{E^ue^{\beta_\pi V}} - \ln(E^ue^{\beta_\pi V})\right) + \frac{\kappa_\rho}{\Delta}\left[\rho(\beta_\pi)\Delta\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+\Delta} - \ln\left(1-\bar{\rho}\Delta + \bar{\rho}\Delta e^{\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+\Delta}}\right)\right]\\
&= \kappa_\pi\beta_\pi\rho(\beta_\pi)\left(\frac{E^uVe^{\beta_\pi V}}{E^ue^{\beta_\pi V}}\right) - \kappa_\pi\beta_\pi\rho(\beta_\pi)W_{t+\Delta} - \frac{\kappa_\rho}{\Delta}\ln\left(1-\bar{\rho}\Delta + \bar{\rho}\Delta e^{\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+\Delta}}\right)\\
&= \kappa_\pi\beta_\pi\rho(\beta_\pi)\left[\frac{E^uVe^{\beta_\pi V}}{E^ue^{\beta_\pi V}} - W_{t+\Delta}\right] - \frac{\kappa_\rho}{\Delta}\ln\left(1-\bar{\rho}\Delta + \bar{\rho}\Delta e^{\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+\Delta}}\right) \;\equiv\; t(\beta_\pi), \qquad (97)
\end{aligned}
$$

where $\rho(\beta_\pi)$ is the function defined by (81). Note that, conditional on $\vec{V}_{t+\Delta}$ and $W_{t+\Delta}$, the right-hand side of (97) can be viewed as a function of β_π, which we call $t(\beta_\pi)$; it represents the total time devoted to decision-making, as a function of precision β_π.

Under very weak assumptions, we can show that $t(\beta)$ slopes upward from $t(0) = 0$ at zero precision. Therefore the curve $h = 1 - t(\beta)$ given by (95) slopes downward, and hence the model has a unique interior optimum.

Proposition 8 *Suppose $f'(h) > 0$ and $f''(h) \leq 0$ for $h \in [0, 1]$, and $\min_i V_i < \max_i V_i$. Then problem (60) is solved by the pair $h^* \in (0, 1)$, $\beta^*_\pi > 0$ that correspond to the unique crossing of curves (95)-(96). The solution (h^*, β^*_π) varies smoothly with changes in $\vec{V}$ and $W_{t+\Delta}$.*

Proof. The main point to prove is that the right-hand side of (97) is nondecreasing in β. Let us write (97) as $h = 1 - t(\beta)$, where

$$t(\beta) \;=\; \kappa_\pi \beta \rho(\beta) \left(\frac{m'(\beta)}{m(\beta)} - W \right) - \frac{\kappa_\rho}{\Delta} \ln \left(1 - \bar{\rho}\Delta + \bar{\rho}\Delta \left(\frac{m(\beta)}{e^{\beta W}} \right)^{\frac{\kappa_\pi}{\kappa_\rho}} \right), \tag{98}$$

$m(\beta) \equiv E^u \exp(\beta V)$, and $\rho(\beta) = \bar{\rho} \exp\left(\frac{\kappa_\pi}{\kappa_\rho} \beta_\pi D_{t+\Delta} \right)$.[17] Note that $\rho(\beta) \in [0, 1]$ for $\beta \in [0, \infty)$. Also, $\frac{m'(0)}{m(0)} = E^u V$, and $m(0) = 1$. Therefore $t(0) = 0$.

Now by differentiating and simplifying, we can show that

$$\begin{aligned} t'(\beta) \;&=\; \kappa_\pi \rho \left(\frac{m'}{m} - W \right) + \kappa_\pi \beta \rho' \left(\frac{m'}{m} - W \right) + \kappa_\pi \beta \rho \left(\frac{m''}{m} - \left(\frac{m'}{m} \right)^2 \right) \\ &\quad - \frac{\kappa_\rho}{\Delta} \frac{\kappa_\pi}{\kappa_\rho} \frac{\bar{\rho}\Delta \left(\frac{m(\beta)}{e^{\beta W}} \right)^{\frac{\kappa_\pi}{\kappa_\rho} - 1}}{1 - \bar{\rho}\Delta + \bar{\rho}\Delta \left(\frac{m(\beta)}{e^{\beta W}} \right)^{\frac{\kappa_\pi}{\kappa_\rho}}} \left(\frac{e^{\beta W} m' - m W e^{\beta W}}{e^{2\beta W}} \right) \\ &=\; \kappa_\pi \rho \left(\frac{m'}{m} - W \right) + \kappa_\pi \beta \rho' \left(\frac{m'}{m} - W \right) + \kappa_\pi \beta \rho \left(\frac{m''}{m} - \left(\frac{m'}{m} \right)^2 \right) - \kappa_\pi \rho \left(\frac{m'}{m} - W \right) \end{aligned}$$

Notice that the first and last terms cancel; the remaining terms are

$$t'(\beta) \;=\; \frac{\kappa_\pi^2}{\kappa_\rho} \rho(\beta) \left(\frac{m'}{m} - W \right)^2 + \kappa_\pi \beta \rho \left(\frac{m''}{m} - \left(\frac{m'}{m} \right)^2 \right). \tag{99}$$

The first term on the right-hand side of (99) is obviously nonnegative. In the second term on the right-hand side of (99), we have

$$\frac{m''}{m} - \left(\frac{m'}{m} \right)^2 = \frac{E^u V^2 \exp(\beta V)}{E^u \exp(\beta V)} - \left(\frac{E^u V \exp(\beta V)}{E^u \exp(\beta V)} \right)^2 = E^\pi V^2 - (E^\pi V)^2 \geq 0. \tag{100}$$

[17] In the continuous-time limit, this reduces to

$$t(\beta) \;=\; \kappa_\pi \beta \rho(\beta) \left(\frac{m'(\beta)}{m(\beta)} - W \right) + \kappa_\rho (\bar{\rho} - \rho(\beta)),$$

which simplifies the derivative:

$$t'(\beta) \;=\; \kappa_\pi \rho \left(\frac{m'}{m} - W \right) + \kappa_\pi \beta \rho' \left(\frac{m'}{m} - W \right) + \kappa_\pi \beta \rho \left(\frac{m''}{m} - \left(\frac{m'}{m} \right)^2 \right) - \kappa_\rho \rho'.$$

The first and last terms cancel, to give (99).

This is the quantity that we signed in Lemma 1(c); since it represents a variance, it is *strictly* positive for all $\beta \in [0, \infty)$ as long as $\min_i V_i < \max_i V_i$. Therefore both terms on the right-hand side of (99) are nonnegative, and in particular, we have $t'(\beta) > 0$ strictly for all $\beta \in (0, \infty)$ under the maintained assumptions of the proposition.

Thus, the curve $h = 1 - t(\beta_\pi)$ slopes down from $h = 1$ at $\beta_\pi = 0$. The curve $f'(h) = \frac{1-\delta\Delta}{\kappa_\pi \beta_\pi}$ slopes upward from $\beta_0 \equiv \frac{1-\delta\Delta}{\kappa_\pi f'(0)} \geq 0$; it eventually exceeds one if $f'(1) > 0$. This proves that the two curves have a interior unique crossing, as stated in the proposition. **QED.**

Thus we can also characterize the comparative statics.

$$
\begin{aligned}
-dh &= \frac{\partial t}{\partial \beta_\pi} d\beta_\pi + \frac{\partial t}{\partial W} dW_{t+\Delta} + \left(\frac{\partial t}{\partial V}\right)' d\vec{V}_{t+\Delta} && (101)\\
\kappa_\pi \frac{\partial f}{\partial h} d\beta_\pi + \kappa_\pi \beta_\pi \frac{\partial^2 f}{\partial h^2} dh + \kappa_\pi \beta_\pi \frac{\partial^2 f}{\partial A \partial h} &= 0 && (102)
\end{aligned}
$$

which can be summarized as

$$
\begin{pmatrix} dh \\ d\beta_\pi \end{pmatrix} = \frac{1}{|J|} \begin{pmatrix} -\kappa_\pi \frac{\partial f}{\partial h} \frac{\partial t}{\partial W} & -\kappa_\pi \frac{\partial f}{\partial h} \left(\frac{\partial t}{\partial V}\right)' & -\frac{\partial t}{\partial \beta_\pi} \kappa_\pi \beta_\pi \frac{\partial^2}{\partial A \partial h} \\ \kappa_\pi \beta_\pi \frac{\partial^2 f}{\partial h^2} \frac{\partial t}{\partial W} & \kappa_\pi \beta_\pi \frac{\partial^2 f}{\partial h^2} \left(\frac{\partial t}{\partial V}\right)' & \kappa \pi \beta_\pi \frac{\partial^2}{\partial A \partial h} \end{pmatrix} \begin{pmatrix} dW_{t+\Delta} \\ d\vec{V}_{t+\Delta} \\ dA \end{pmatrix} \qquad (103)
$$

where

$$
J = -\kappa_\pi \frac{\partial f}{\partial h} + \frac{\partial t}{\partial \beta_\pi} \kappa_\pi \beta_\pi \frac{\partial^2 f}{\partial h^2} < 0.
$$

2.4.2 Numerical example

As we saw in the previous section, assuming time is valuable is very helpful, since it ensures an interior solution for the precision of the decision problem. Since there is only one unit of time available, the highest possible precision consistent with an interior solution is $\beta_\pi = \beta_h^* \equiv \frac{1-\delta\Delta}{\kappa_\pi} f'(1)$. Assuming that $\lim_{h \to 0} f'(h) = \infty$, the lowest possible precision is zero. Hence we can limit the search for the optimal level of precision β_π to the finite set $[0, \beta_1^*]$.

Moreover, given a hypothetical precision level β_π, the first-order conditions of (60) imply that the precision on the timing decision must be $\beta_\rho = \frac{\kappa_\pi}{\kappa_\rho} \beta_\pi$, and the fraction of time spent working must be $h = (f')^{-1} \left(\frac{1-\delta\Delta}{\kappa_\pi \beta_\pi}\right)$. Finally, we can use (72), (76), (**??**), and (85) to calculate $\mathcal{D}_\pi$, $D_{t+\Delta}$, ρ, $\mathcal{D}_\rho$, so we can calculate the amounts of time $\tau = \kappa_\pi \rho \mathcal{D}_\pi$ and $\mu = \kappa_\rho \mathcal{D}_\rho$ dedicated to the two aspects of the decision problem. So solving for the optimal decision only requires finding a β_π consistent with $h + \tau + \mu = 1$.

Figures 5-6 illustrate the behavior of this model of decision with two numerical examples. We choose some arbitrary parameters in order to be able to solve the problem: the time step is $\Delta = 0.005$; the discount factor is $\delta = 0.01$; the information cost parameters are $\kappa_\pi = \kappa_\rho = 0.1$; the production function is $f(h) = Ah^\alpha$, with $A = 2$ and $\alpha = 2/3$. The continuation value of the problem is $W_{t+\Delta} = 10$. We consider two possible configurations of the decision options. First, in Fig. 5, we assume the options have low value:

$$
\vec{V}_{low} \equiv (V_1, V_2, \ldots, V_{13}, V_{14}) \equiv (-20, -18, \ldots, 4, 6).
$$

In Fig. 6, we add five to all the payoffs, so that the options are more valuable:

$$
\vec{V}_{high} \equiv (V_1, V_2, \ldots, V_{13}, V_{14}) \equiv (-15, -13, \ldots, 9, 11).
$$

Figure 5: Time use in equilibrium: example with low values.

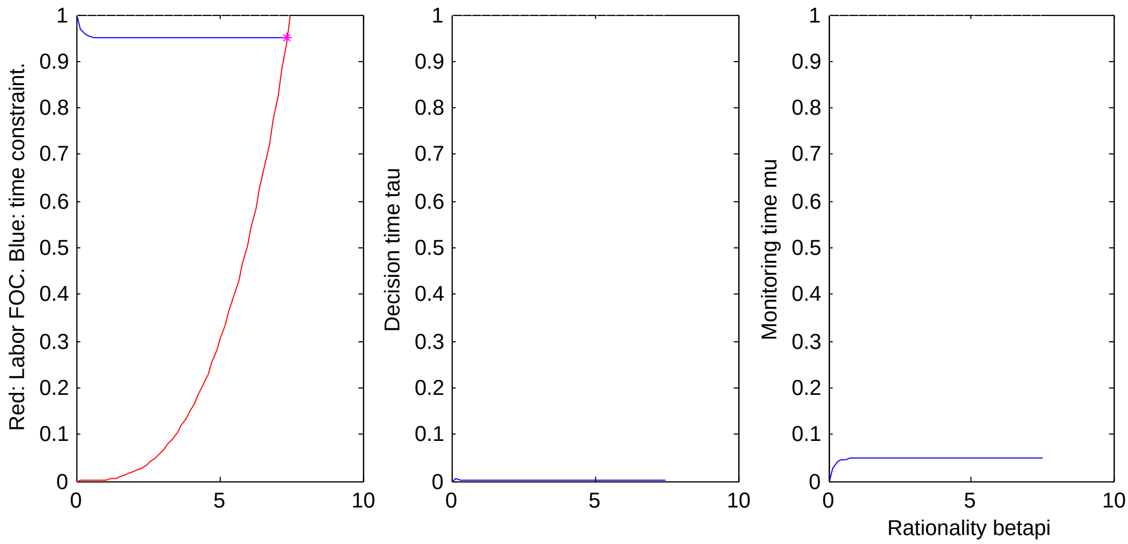

Note: Time use in problem with low values, $\max_i V_i < W_{t+\Delta}$.
Left panel: First-order condition (96) in red, and time constraint $h = 1 - \tau - \mu$ in blue.
Middle and right panels: fractions of time τ and μ devoted to decision and monitoring, respectively.

The figures show β_π on the horizontal axes of all the panels. The relevant range of precisions is $\beta_\pi \in [0, \beta_1^*] \approx [0, 7.5]$. The first panel of each figure shows the two equations that determine equilibrium precision β_π^* (highlighted with a red star). The red line shows the first-order condition (96) that equates the marginal payoff of time devoted to the decision with the marginal value of other uses of time; this gives h as an increasing function of β_π. The blue line shows the time constraint (95), representing the time left over (h) after the time devoted to the decision (as a function of β_π). The next two panels break the decision time into its two component parts: the time τ required to choose across the various options with precision β_π, and the time μ required to control the timing of the decision with accuracy $\beta_\rho = \frac{\kappa_\rho}{\kappa_\pi}\beta_\pi$.

Note that given the low payoffs assumed in the example of Figure 5, it would be optimal to simply postpone the decision; even choosing the best option (V_{14}) with probability one, the value of postponement is higher. But the imperfect decision maker assumed here is not entirely sure of this, and can only draw a conclusion be devoting some time to thinking about it. Thus we see in the third panel that raising μ to approximately 0.05 suffices to achieve arbitrarily high timing accuracy β_ρ. As this accuracy increases, the DM is increasingly convinced that the decision should simply be postponed (and accordingly decreases the arrival rate from its benchmark level $\bar{\rho} = 0.5$ towards zero); therefore the time devoted to the decision itself remains almost exactly at zero. Thus, as the red star in the first panel indicates, in this example approximately 95% of time is devoted to non-decision activities, 5% of time is devoted to learning that the decision is not worthwhile, and the probability of actually completing the decision is extremely close to zero.

Figure 6: Time use in equilibrium: example with high values.

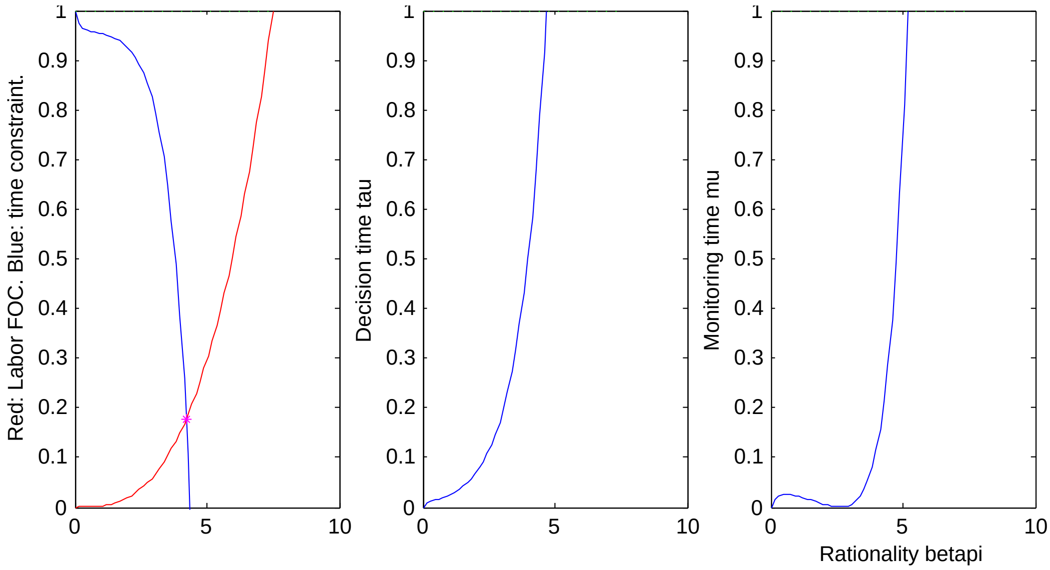

Note: Time use in problem with one high value, $\max_i V_i > W_{t+\Delta}$.
Left panel: First-order condition (96) in red, and time constraint $h = 1 - \tau - \mu$ in blue.
Middle and right panels: fractions of time τ and μ devoted to decision and monitoring, respectively.

On the other hand, in Fig. 6, the best option $\max_i V_i = 11$ slightly exceeds $W_{t+\Delta} = 10$; therefore, making the decision is potentially welfare improving, but only if precision is sufficiently high. In the second panel of the figure, we see that increasing levels of β_π are accompanied by steadily increasing time τ devoted to the decision. In the third panel, for very low β_π, a small amount of time is devoted to evaluating whether the decision should be made; the answer is *no*, conditional on this low accuracy, so the decision-maker decreases the arrival rate below the benchmark rate (setting $\rho < \bar{\rho} = 0.5$) over this range. But at sufficiently high precision, roughly $\beta_\pi > 3$, the probability of making the right decision becomes high enough that completing the decision is worthwhile. Therefore for $\beta_\pi > 3$, the monitoring time μ steadily increases, implying in this case that the DM is raising the arrival rate above its benchmark rate ($\rho > \bar{\rho} = 0.5$). Note however that at roughly $\beta_\pi > 3$, completing and postponing the decision are equally valuable ($\tilde{V}_{t+\Delta} = W_{t+\Delta}$), so the DM devotes no time to evaluating whether or not the decision should be completed. Therefore the DM sets $\mu = 0$, implying $\rho = \bar{\rho}$ exactly.

This nonmonotonicity of time devoted to monitoring whether completing the decision is worthwhile explains why the blue time constraint curve in the first panel of Fig. 6 has an inflection point, going from convex to concave. Nonetheless, as we showed in Proposition 8, total decision time as a function of precision β_π is unambiguously increasing. Therefore the blue curve slopes unambiguously downwards, and hence a unique crossing of the two curves exists, at an interior point. In the example of Fig. 6, more than 80% of time available is devoted to completing the decision; most of this is actual decision time τ rather than monitoring time μ.

3 Bargaining games

We now apply our decision framework to some simple sequential bargaining games. In the games, two impatient players, A and B, negotiate shares of a cake. Each may propose a share of the cake between 0 and 1 to the other player, and may accept or reject an offer received from the other player. Accepting the offer of the other player ends the game. We allow for costly, error-prone decisions when making offers, and when deciding whether or not to accept them, so these games will be applications of the decision framework described in Prop. 8.

We compare two assumptions about the bargaining protocol:

- One particularly simple protocol is an **alternating offers protocol** where "reject" is equated with making a different proposal, and the proposer of the current offer *cannot revise* that offer. Thus this protocol has only two types of decisions: making an offer, and accepting or rejecting an offer.

- A simple alternative would be a **rejection-first protocol** where offers are rejected by simply saying "no", which returns the game to a situation with no offer outstanding. Again, the most recent proposer *cannot revise* the offer.

Two simple but interesting alternatives would be:

- Another simple extension would be an **updateable offers protocol** where "reject" is equated with making a different proposal, but the proposer of the current offer *can update* that offer with a new one.

- Finally, we could consider a **withdrawable offers protocol** that allows either player to say "no" to the current offer, returning the game to a situation with no offer outstanding. That is, the recipient of the offer can accept it or reject it without offering an alternative; the player that made the offer can likewise withdraw it without offering an alternative.

Note that under the alternating offers protocol, the responder must be careful not to propose an unfavorable alternative, and can take plenty of time to respond since the proposer can do nothing else until the responder acts. But in turn this gives both players an incentive to be very careful when making the first offer. So this protocol could imply substantial delay, compared with the other protocols. The updateable and withdrawable protocols seem like especially robust versions of the game, since they never force players to remain in an undesirable state of the game.

We will solve the games by backwards induction from a finite ending time. Time steps are assumed finite but very short; therefore we ignore the possibility that offers arrive simultaneously. Thus under all the protocols considered, any given player $i \in \{A, B\}$ may be in one of three types of states at any time prior to the end of the game:

- State N_t^i ("none"): No offer outstanding

- State $M_t^i(s)$ ("mine"): Own offer outstanding

- State $R_t^i(s)$ ("received"): Other's offer outstanding

By definition, the state of player A is N_t^A if and only if the state of player B is N_t^B. The state of player i is $M_t^i(s)$ if and only if the state of the other player, called $-i$, is $R_t^{-i}(s)$. To

Figure 7: Node N_t^i: no offer outstanding.

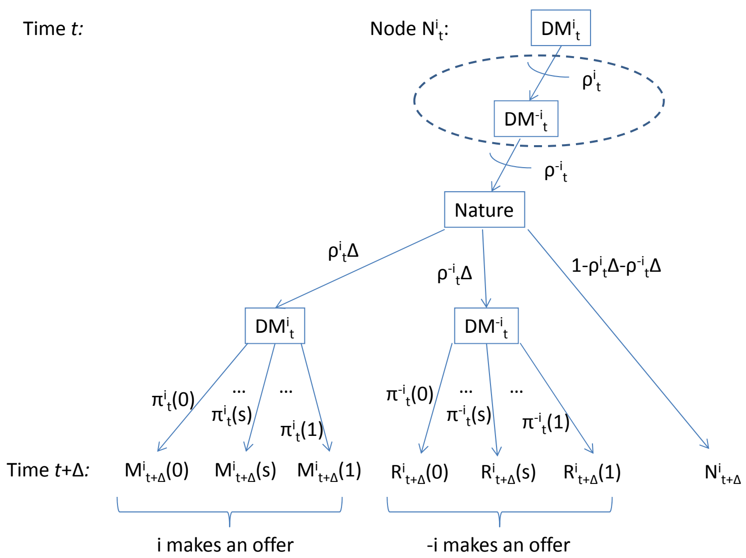

Note: Decision makers A and B simultaneously choose arrival rates ρ_t^A and ρ_t^B of their offers. A slower decision implies higher precision in the allocation of probabilities $\pi_t^i(s)$ across possible offers $s \in [0,1]$.

avoid unnecessary notation, we will identify the names of the states with the names of the value functions associated with those states.

Without loss of generality, we define all offers from the point of view of agent A. Thus $s \in \Gamma^s \subset [0,1]$ represents A's share of the cake. Here Γ^s represents a finite grid of possible shares between zero and one. Terminal payoffs are $u^A(s) \equiv s$ for player A, and $u^B(s) \equiv 1-s$ for player B.

As we discussed in Section 2.1, the precisions of players' decisions are assumed to be bounded by expected time spent thinking about the decisions. Let ρ_t^i indicate the arrival rate of an offer by player i at time t, and let $\pi_t^i(s)$ indicate the probabilities of the possible offers s associated with an offer made by i at t; then the precision of the distribution π_t^i is constrained by τ_t^i/ρ_t^i, where τ_t^i is the fraction of time i devotes to thinking about the decision. Similarly, let $\lambda_t^i(s)$ indicate the arrival rate of an acceptance or rejection by player i at time t; let $\alpha_t^i(s)$ and $1-\alpha_t^i(s)$ indicate the acceptance and rejection probabilities for offer s (conditional on arrival of the decision).

The structure of the game can be described by drawing game trees representing the possible events at nodes of type N_t^i, $M_t^i(s)$, and $R_t^i(s)$. Figures 7-9 describe the game associated with the "alternating offers protocol" mentioned above. In Fig. 7, we see that players i and $-i$ simultaneously choose the arrival rates ρ_t^i and ρ_t^{-i} of their decisions at node N_t^i (the figure is drawn from the point of view of player i; it could be redrawn from the perspective of player $-i$ by simply relabelling nodes with the names used from the perspective of that player).[18] Next, Nature chooses to complete player i's decision with probability $\rho_t^i\Delta$, or player $-i$'s decision with

[18]The simultaneous decision is indicated by the dashed oval surrounding the set of nodes DM_t^{-i}, which represents an information set. That is, all instances of the node DM_t^{-i} lie in the same information set, meaning that player $-i$ is unaware which value of ρ_t^i was chosen by player i.

Figure 8: Node $M_t^i(s)$: offer of player i outstanding.

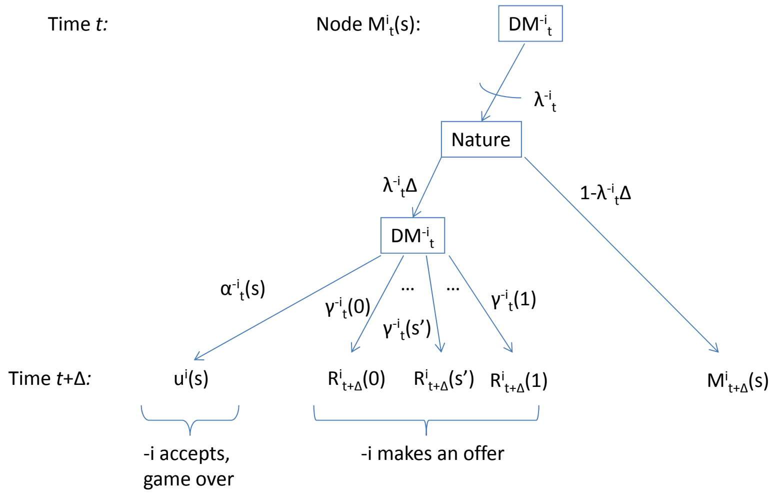

Note: Player $-i$ chooses response arrival rate λ_t^{-i}. A slower response implies higher precision in the decision to accept or to allocate probabilities $\gamma_t^{-i}(s')$ across alternative offers $s' \in [0,1]$.

probability $\rho_t^{-i}\Delta$, or otherwise to go on to node $N_{t+\Delta}^i$ with no offer outstanding at time $t+\Delta$.[19] If player i completes a decision, the game is distributed across nodes $M_{t+\Delta}^i(s)$ at time $t+\Delta$, that is, an offer s made by player i will be on the table. If player i completes a decision, the game will be distributed across nodes $M_{t+\Delta}^i(s)$ at time $t+\Delta$, that is, an offer s made by player i will be on the table. If instead player $-i$ completes a decision, the game will move to one of the nodes $R_{t+\Delta}^i(s)$, indicating that an offer s has been made by player $-i$.

In Figure 8, showing node $M_t^i(s)$, where player i has made an offer s, we see that player $-i$ is choosing a response arrival rate $\lambda_t^{-i} \in [0, \Delta^{-1}]$. Nature permits the response to arrive with probability $\lambda_t^{-i}\Delta$, or otherwise continues to node $M_{t+\Delta}^i(s)$ at time $t+\Delta$. The response may be acceptance, with probability $\alpha_t^{-i}(s)$, which ends the game, so that player i receives payoff $u^i(s)$. Otherwise, player $-i$ may propose an alternative offer s' with probability $\gamma_t^{-i}(s')$, where $\sum_{s'} \gamma_t^{-i}(s') = 1 - \alpha_t^{-i}(s)$; thus player i arrives to node $R_{t+\Delta}^i(s')$ indicating that i must consider offer s' from $-i$. Figure 9, showing a node $R_t^i(s)$ where i has received an offer from $-i$, has exactly the same structure as node $M_t^i(s)$, with appropriate relabelling.

Before going further, it is helpful to discuss why player i's decision at node $R_t^i(s)$ is well-behaved. In Section 2, we considered an arbitrary set of discrete alternatives. But at node $R_t^i(s)$, there are two distinct classes of alternatives: the discrete option to accept, ending the game, or many possible alternative shares $s \in [0,1]$. Numerically, the set of alternative offers is represented as a discrete, uniformly-spaced grid from 0 to 1. But we should ask whether the results might change as we change the density of grid points— or equivalently, is this discrete representation a valid approximation to a model with a continuum of possible shares s on the unit interval?

[19]Since this description of the probability ignores the possibility of simultaneous arrivals, the time step Δ should be chosen short enough so that the implied errors are small.

Figure 9: Node $R_t^i(s)$: offer of player $-i$ outstanding.

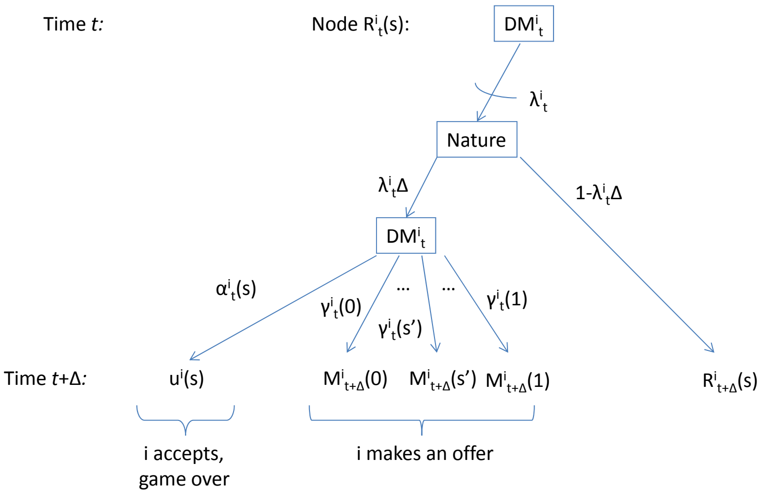

Note: Player i chooses response arrival rate λ_t^i. A slower response implies higher precision in the decision to accept or to allocate probabilities $\gamma_t^i(s')$ across alternative offers $s' \in [0,1]$.

The answer is yes, because of two invariance properties of the relative entropy function. First, under our assumed cost function, choice across a large subset of alternatives Γ has exactly the same cost as choosing initially one subset from a partition of Γ, followed by choosing an element within the chosen subset. Thus we can think of the decision at node $R_t^i(s)$ as two-stage choice, first between acceptance and rejection, and then (immediately) between a set of possible counteroffers. Second, the relative entropy between two distributions is continuous as we go from a discrete support to a continuous support. Thus, as long as we assume a sufficiently fine grid of possible counteroffers s, the results will accurately approximate choice from a continuum of counteroffers $s \in [0,1]$.

3.1 Solving the alternating offers game

We now solve the game defined in Figures 7-9 by backwards induction. We begin from a hypothetical end time T, when further negotiation is impossible. Now suppose we know the solution of the game at time $t+\Delta$, and wish to solve it at time t. It is helpful to consider first the value $R_t^i(s)$, described in Fig. 9, of accepting an offer s made by the other player, or counterproposing. This figure is exactly analogous to Figure 1; it describes a single decision made by a single player subject to an outside option $R_{t+\Delta}^i$ which is already known, since we assume the solution at time $t+\Delta$ is already known.

Thus, suppose there are $n-1$ points $(s_1, \ldots, s_{n-1})$ in the grid of possible offers Γ^s. Given some specific current offer s, consider a vector of possible alternatives

$$(V_1, \ldots, V_n) \equiv (M_{t+\Delta}^i(s_1), \ldots, M_{t+\Delta}^i(s_{n-1}), u^i(s)), \tag{104}$$

chosen with probabilities

$$(\pi_1, \ldots, \pi_n) \equiv (\gamma_t^i(s_1), \ldots, \gamma_t^i(s_{n-1}), \alpha_t^i(s)). \tag{105}$$

Define the continuation value

$$W_{t+\Delta}^i \equiv R_{t+\Delta}^i(s). \tag{106}$$

Suppose that player i has discount rate δ^i and allocates fraction μ of time evaluating whether the problem should be solved, fraction τ to actually solving the problem, and obtains payoff $f(h)\Delta$ from the remaining time $h = 1 - \mu - \tau$. The player faces the information constraints

$$\kappa_\pi^i \left(\sum_{j=1}^{n} \pi_j \ln \pi_j + \ln n \right) \leq \frac{\tau}{\lambda_t^i}, \tag{107}$$

where λ_t^i is the chosen arrival rate of the solution, and

$$\kappa_\lambda^i \left(\lambda_t^i \Delta \ln \left(\frac{\lambda_t^i}{\bar{\lambda}^i} \right) + (1 - \lambda_t^i \Delta) \ln \left(\frac{1 - \lambda_t^i \Delta}{1 - \bar{\lambda}^i \Delta} \right) \right) \leq \mu \Delta. \tag{108}$$

Here we have defined a decision problem of the form analyzed in Proposition 8. To ensure that the solution is interior, we assume that time has some value outside of the bargaining game: $f'(h) > 0$ and $f''(h) \leq 0$. Then the proposition implies that the rationality coefficient β and the time in alternative activities h are given by the unique crossing of the two curves (95)-(96).[20] Given β and h, we can also calculate τ, μ, λ_t^i, and the logit probabilities $(\gamma_t^i(s_1), \ldots, \gamma_t^i(s_{n-1}), \alpha_t^i(s))$. Once we know h, λ, and the logit probabilities, we can update the value function as follows:

$$R_t^i(s) = f(h)\Delta + (1 - \delta^i \Delta) \left\{ (1 - \lambda_t^i \Delta) R_{t+\Delta}^i + \lambda_t^i \Delta \left(\sum_{j=1}^{n-1} \gamma_t^i(s_j) M_{t+\Delta}^i(s_j) + \alpha_t^i(s) u^i(s) \right) \right\}. \tag{109}$$

Next, consider the value function $M_t^i(s)$ defined in Figure 8. Note that the only decisions made in Fig. 8 are decisions of the *other* agent $-i$. Therefore, agent i can set $h = 1$, dedicating no time to decisions. And thus after we have solved the other agent's problem to calculate $R_t^{-i}(s)$ and the associated logit probabilities and arrival rate, we can calculate i's value $M_t^i(s)$ by summing up the possible events shown in the game tree:

$$M_t^i(s) = f(1)\Delta + (1 - \delta^i \Delta) \left\{ (1 - \lambda_t^{-i} \Delta) M_{t+\Delta}^i + \lambda_t^{-i} \Delta \left(\sum_{j=1}^{n-1} \gamma_t^{-i}(s_j) R_{t+\Delta}^i(s_j) + \alpha_t^{-i}(s) u^i(s) \right) \right\}. \tag{110}$$

The game at a node where no offer has yet been made is somewhat more complicated, because, even taking as given the time $t + \Delta$ equilibrium, the optimal decision of player A at time t interacts with that of player B at t. Nonetheless, the problem is tractable because the

[20] A MATLAB program, twomargins_ns.m, is provided to solve the problem analyzed in Proposition 8. That program is called repeatedly by the program rubegame_alternating.m to solve each backwards induction step, in order to calculate the equilibria reported here.

interaction goes *only* through the continuation values $W^A_{t+\Delta}$ and $W^B_{t+\Delta}$. Therefore each player's decision can be written in the form considered in Proposition 8.

At node N^i_t, player i chooses between options

$$(V_i, \ldots, V_{n-1}) \;\equiv\; (M^i_{t+\Delta}(s_1), \ldots, M^i_{t+\Delta}(s_{n-1})), \tag{111}$$

chosen with probabilities

$$(\pi_1, \ldots, \pi_{n-1}) \;\equiv\; (\pi^i_t(s_1), \ldots, \pi^i_t(s_{n-1})). \tag{112}$$

We write the fractions of time dedicated to monitoring, choosing, and alternative activities as μ, τ, and h, respectively; and the arrival rate chosen by i as ρ^i_t. The constraints on the decision are

$$\frac{\tau}{\rho^i_t} \geq \kappa^i_\pi \left(\sum_{j=1}^{n} \pi_j \ln \pi_j + \ln(n-1) \right), \tag{113}$$

$$\mu\Delta \geq \kappa^i_\rho \left(\rho^i_t \Delta \ln\left(\frac{\rho^i_t}{\bar{\rho}^i}\right) + (1-\rho^i_t\Delta)\ln\left(\frac{1-\rho^i_t\Delta}{1-\bar{\rho}^i\Delta}\right)\right). \tag{114}$$

Now, if player i fails to make a decision at node N^i_t, he will face an offer from $-i$ with probability $\rho^{-i}_t\Delta$, or will return to a node with no offer outstanding with probability $1-\rho^{-i}_t\Delta$. Therefore the continuation value for player i's problem is

$$W^i_{t+\Delta} \;=\; (1-\rho^{-i}_t\Delta)N^i_{t+\Delta} + \rho^{-i}_t\Delta \sum_{j=1}^{n-1} \pi^{-i}_t(s_j) R^i_{t+\Delta}(s_j). \tag{115}$$

With this notation we have defined a decision of the form analyzed in Prop. 8, but it cannot be solved directly because $W^i_{t+\Delta}$ is unknown: it depends on the as-yet-unknown solution to player $-i$'s problem. Thus, we can instead make the initial guess $W^i_{t+\Delta} \approx N^i_{t+\Delta}$. We can then solve the problem, according to Prop. 8, and calculate its value, as given by

$$N^i_t = f(h)\Delta + (1-\delta^i\Delta)\left\{ (1-\rho^i_t\Delta)W^i_{t+\Delta} + \rho^i_t\Delta \sum_{j=1}^{n-1} \pi^i_t(s_j) M^i_{t+\Delta}(s_j) \right\} \tag{116}$$

$$= f(h)\Delta + (1-\delta^i\Delta)\left\{ (1-\rho^i_t\Delta)W^i_{t+\Delta} + \rho^i_t\Delta \ln\left(\frac{1}{n-1}\sum_{j=1}^{n-1} \exp(\beta M^i_{t+\Delta}(s_j)) \right) \right\}. \tag{117}$$

We perform the same calculations for $-i$ that we have performed for i. We can then update our guesses for $W^i_{t+\Delta}$ and $W^{-i}_{t+\Delta}$ using the formula (115), and solve both players' problems again. We continue iterating in this way until we find fixed points for $W^i_{t+\Delta}$, N^i_t, $W^{-i}_{t+\Delta}$, and N^{-i}_t.

We have now described a single backwards induction step to update the values $R^i_t(s)$, $M^i_t(s)$, and N^i_t, and likewise for $-i$. Working backwards until all value functions have converged, we arrive at an equilibrium of the game. Figure 10 illustrates a numerical solution of this equilibrium, for two risk-neutral agents that split a cake of size 100, under the parameters $\Delta = 0.01$, $\kappa = 0.2$, $\delta = 0.05$, and $\bar{\rho} = 0.5$, $A = 1$, and $\alpha = 2/3$, which are chosen for computational convenience. A discrete grid of offers is considered, from $s = 0$ to $s = 100$ by

Figure 10: Equilibrium of alternating offers game.

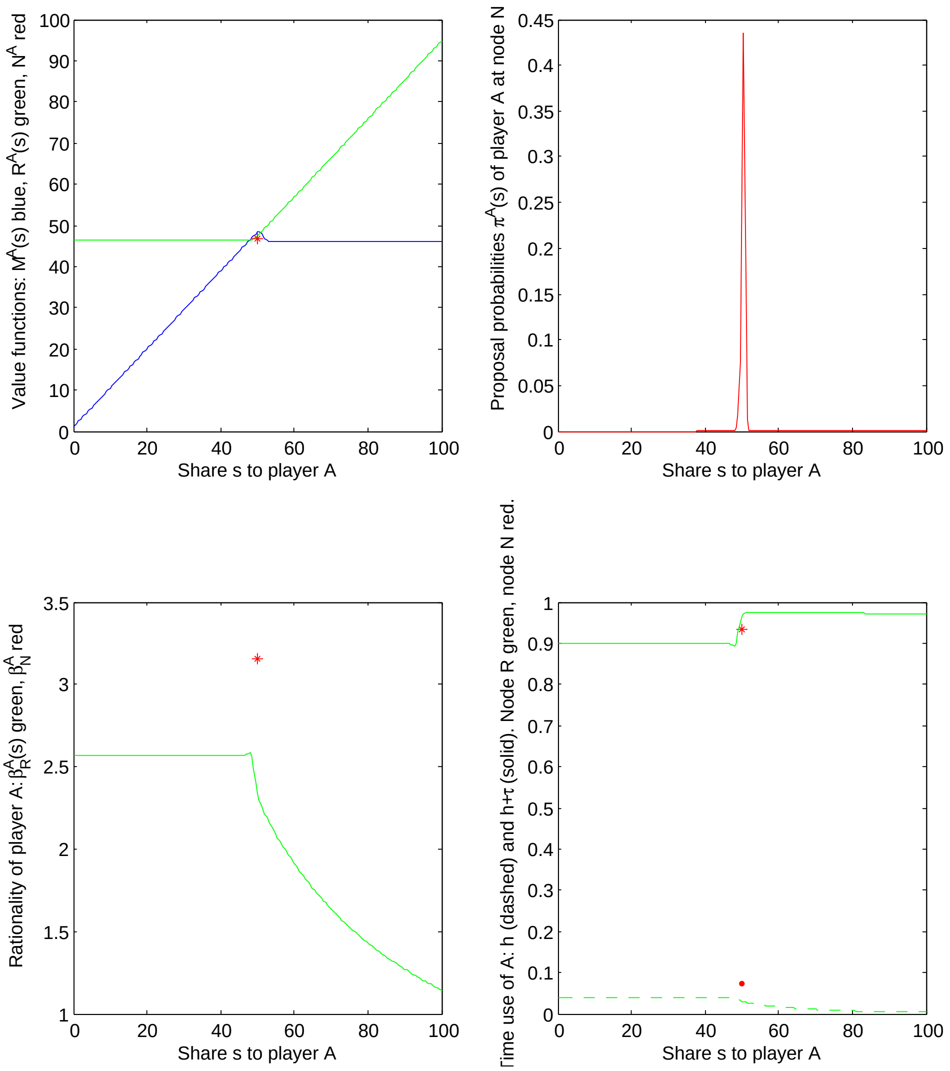

Note: Equilibrium of alternating offers game, including errors in timing (choices described by Prop. 8). Green: Node R; blue: node M; red: node N.

Top left: Values to A of outstanding offers. Blue: Value $M^A(s)$ of offer made by A. Green: Value $R^A(s)$ of offer made by B. Red: Value N^A when no offer has been made.
Top right: Offer probabilities $\pi^A(s)$ when A makes an offer at node N.
Bottom left: Rationality of A. Green: rationality $\beta_R^A(s)$ when responding to offer s. Red: Rationality β_N^A when making an offer at node N.
Bottom right: Time use of A at nodes $R^A(s)$ and N^A. Plus: labor h; star: $h + \tau$.

steps of 0.5. The time step Δ is chosen small enough so that the probability of making a decision in a single time step is close to zero, implying that any given decision will take many time steps. The offer step size is chosen fine enough to guarantee that equilibrium is unique (for further discussion, see Sec. 3.5 below). The discount rate δ is assumed relatively large, to speed up the equilibrium calculation (these examples take two or three hours to compute on an ordinary laptop computer). The precision cost κ is chosen large enough so that errors are visible in the graphs, without displaying a high degree of irrationality. The settings of $\bar{\rho}$ and α are essentially arbitrary. The value of A is a normalization, but its setting relative to the size of the cake is crucial in one sense. We choose parameters so that devoting time to splitting the cake is in fact worthwhile. In a symmetric example, this requires that the present discounted value of never bargaining, and instead devoting all future time to alternative activities, $A/\delta = 20$, is less than the value of immediately ending the game by splitting the cake evenly, which is 50.

The first panel graphs the functions $M^A(s)$ and $R^A(s)$ that represent player A's value of an outstanding offer when A made the offer, or received the offer, respectively. Both are graphed as a function of the share s received by A, so $M^A(s)$ and $R^A(s)$ are mostly constant or increasing in s. This equilibrium is symmetric around an offer of a 50/50 split, so $M^B(1-s) = M^A(s)$ and $R^B(1-s) = R^A(s)$; therefore $M^B(s)$ and $R^B(s)$ (not shown) are mostly constant or decreasing.[21] The value to A of an offer made by A, $M^A(s)$, is shown in blue. It starts at 0 for $s = 0$, and is initially approximately linearly increasing, peaking at $M^A(50.5) = 48.47$. It then falls slightly to a flat plateau at $M^A(s) = 46.20$ for offers in the range $s \geq 54$. Values in this range reflect the fact that player B accepts offers less than 50 with probability near one, as we can see in the first panel Figure 11. Higher offers are would be more valuable to A if B accepted them, but the acceptance probability is less than one percent for $s \geq 53$, so $M^A(s)$ at high values of s mostly reflects the value of waiting for an alternative offer made by B.

While A's value is high conditional on having made a high offer, it is even higher conditional on *receiving* a high offer, since A has the option of accepting, and ending the game. Thus, we see that the green curve $R^A(s)$, which represent's A's value of responding to an offer received from B, lies slightly below the 45^o-line for $s \geq 50$. And when $s \leq 47$, we have $R^A(s) = 46.42$, representing the value of the option to reject. While it might appear on first sight that $R^A(s) \geq M^A(s)$ for all s, this is not actually true around $s \approx 50$. In particular, we have $R^A(50) = 47.48 < M^A(50) = 48.31$. The reason is that regardless of who proposed $s = 50$, A can expect to receive the terminal payoff $u^A(50) = 50$ with very high probability very soon. But if A made the offer, then A only needs to wait in order to receive this terminal payoff; if instead B made the offer, then A needs to *think about it* in order to decide to accept. This is why the value of receiving an offer of an even split is very slightly less than the value of having already made the same offer. Both of these values also exceed $N^A = N^B = 46.78$, which is the value when no offer is yet outstanding, reflecting the fact that players will need to think, and time will elapse, prior to receiving a terminal payoff.

Thus equilibrium outcomes are simple and highly rational in spite of the presence of errors in the model: the actual split is likely to be very close to 50/50. The modal proposal by A is 50.5, and 85% of all offers made by A lie in the interval [50,51] (top, right panel of Fig. 10). An even split is accepted with more than 90% probability, while B accepts the proposal $s = 51$ with 56% probability (first panel of Fig. 11). However, some low-cost errors are also observed. For

[21] This equilibrium is computed by starting with a symmetric guess of the form $N^A_{init} = N^B_{init}$. But as long as the step size in the offer grid is sufficiently small, the same (unique) equilibrium is obtained starting from an asymmetric guess. See Sec. 3.5.

Figure 11: Acceptance decision in alternating offers game.

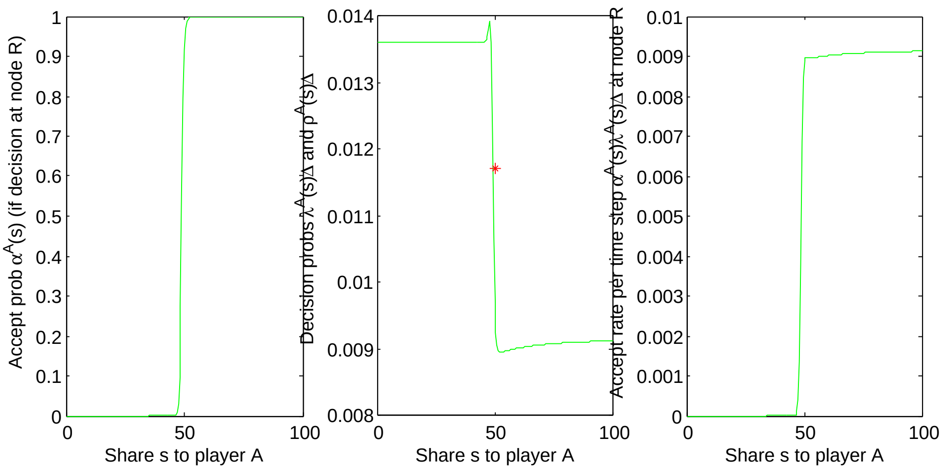

Note: Acceptance probabilities in alternating offers game, including errors in timing (choices described by Prop. 8). Green: Node R; red: node N.

Left panel: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

Middle panel: Decision arrival probabilities of player A: $\lambda^A(s)\Delta$ (green) and $\rho^A(s)\Delta$ (red).

Right panel: Acceptance probabilities $\alpha^A(s)\lambda^A(s)\Delta$ of player A.

example, in the top, right panel of Fig. 10, A has a tiny, roughly constant probability (0.03%) of making any offer s in the range $s > 55$. These unacceptable offers are again relatively low-cost errors, because they simply postpone arrival of an agreement. On the other hand, the probability that A makes an offer below 50 declines extremely rapidly, becoming reaching zero at machine precision for $s \leq 41$. Low offers are obviously much more costly mistakes to A, since they are accepted by B with very high probability.

While behavior in this example is very close to rationality, Figure 11 shows the obvious principal difference from the model of Rubinstein (1982): agreement is not reached immediately. The first panel shows that, conditional on making a decision, the responder accepts any offer $s \geq 50$ with probability α close to one. However, in the second panel, we see that the response to an offer arrives with probability $\lambda\Delta$ near 1% in any single time step. Responses arrive most slowly when the offer is near $s = 50$, since the correct response is least obvious near the acceptance threshold. The probability $\rho\Delta$ of making an offer when no offer is outstanding, per time step, is similar (red star in second panel). Thus, typically, on the order of one hundred time steps pass before an offer is made, and another hundred before it is accepted. The third panel shows the product $\alpha\lambda\Delta$ that represents the probability that any outstanding offer s is accepted in any given time step.

3.2 Solving the game with rejection first

The game we just solved equates rejecting an offer with making a counteroffer. We next consider a game that distinguishes between the two, allowing the responding player to simply say “no”, which returns the game to a node with no offer outstanding. The game tree at state N is the same one we showed in Fig. 7. The games at states $M^i_t(s)$ and $R^i_t(s)$ are changed (as seen in Figs. 8 and 9) by replacing the transitions to a new offer s' by a transition back to state N with probability $1-\alpha^{-i}_t(s)$ or $1-\alpha^i_t(s)$ (respectively). Here $1-\alpha^i(s)$ represents player i's probability of saying "no" to offer s.

Thus the Bellman equation for received offers becomes an application of Prop. 8 with just two options:

$$R^i_t(s) \;=\; f(h)\Delta + (1-\delta^i\Delta)\left\{(1-\lambda^i_t\Delta)R^i_{t+\Delta} + \lambda^i_t\Delta\left(\alpha^i_t(s)u^i(s) + (1-\alpha^i_t(s))N^i_{t+\Delta}\right)\right\}. \quad (118)$$

The value function $M^i_t(s)$ likewise contemplates the possibility that the offer is accepted, with value $u^i(s)$, or rejected, implying a return back to value $N^i_{t+\Delta}$.

$$M^i_t(s) \;=\; f(1)\Delta + (1-\delta^i\Delta)\left\{(1-\lambda^{-i}_t\Delta)M^i_{t+\Delta} + \lambda^{-i}_t\Delta\left(\alpha^{-i}_t(s)u^i(s) + (1-\alpha^{-i}_t(s))N^i_{t+\Delta}\right)\right\}. \quad (119)$$

The Bellman equation for node N^i_t is unchanged from the previous protocol.

The results are shown in Figs. 12-13. Parameters are identical to the previous example (except that we use a finer step size in the offer grid, 0.1, to guarantee equilibrium uniqueness). As in the previous example, under this protocol the offer that is actually accepted remains very close to a 50/50 split, with a slight advantage to the proposer. In this example, the modal offer by A is $s = 51.2$, and the probability that A's offer lies in [49,52] is 93.4%. The main difference in this new specification is that acceptable offers are accepted much more quickly under this rejection-first protocol than they were under the previous one; the probability of accepting an offer of a 50/50 split rises to rises to roughly 3% per time step, from roughly 1% per time step before (comparing the green lines in second panels of Figs. 11 and 13). The reason is intuitive: rejection was a more complex decision under the previous protocol, where rejection implied proposing an alternative. Therefore, the acceptance/rejection decision was much slower under the "alternating offers" protocol than it is under the "reject first" protocol. Note in contrast that the arrival rate of the initial offer (red stars in Figs. 11 and 13) is roughly similar (around 1.2% per time step) under the two specifications, since in both cases the initial proposal requires a choice across the entire offer space.

Thus, the equilibrium under this protocol is a Pareto improvement relative to that arising under the alternating-offers protocol considered previously. The actual allocations are very similar (a split near 50/50 with probability near one), but there is less bargaining delay; therefore the *ex ante* value of the game rises compared with the previous protocol, to $N^A = N^B = 48.18$. Since this protocol is preferred by both players, from here on we will take the reject-first game as our benchmark: all comparative statics calculations and other specification changes will be constructed as deviations away from the equilibrium displayed in Figs. 12-13.

3.3 Games without errors in timing

We now go on to further characterize the behavior of the game, and to study the robustness of our previous results, by considering some variations on the specifications shown so far. From here on, figures for additional specifications studied are moved to the appendix.

First, we consider the role of timing errors. Thus far, we have computed bargaining equilibria in which players were subject to errors both in their choices across available actions, and in the timing of their choices (that is, we applied the specification of Prop. 8, from Section 2.4). Now we

Figure 12: Equilibrium of rejection-first game.

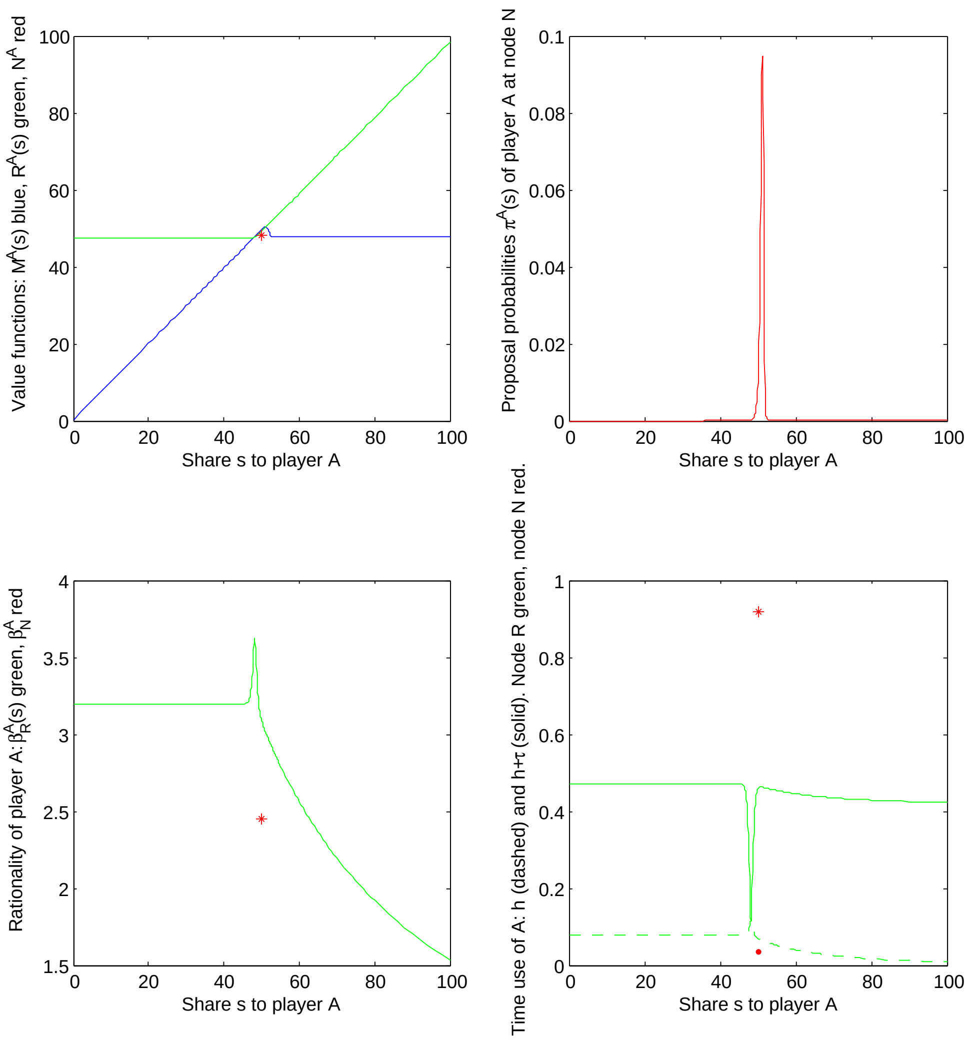

Note: Equilibrium of rejection-first game, including errors in timing (choices described by Prop. 8). Green: Node R; blue: node M; red: node N.

Top left: Values to A of outstanding offers. Blue: Value $M^A(s)$ of offer made by A. Green: Value $R^A(s)$ of offer made by B. Red: Value N^A when no offer has been made.
Top right: Offer probabilities $\pi^A(s)$ when A makes an offer at node N.
Bottom left: Rationality of A. Green: rationality $\beta^A_R(s)$ when responding to offer s. Red: Rationality β^A_N when making an offer at node N.
Bottom right: Time use of A at nodes $R^A(s)$ and N^A. Plus: labor h; star: $h + \tau$.

Figure 13: Acceptance decision in rejection-first game.

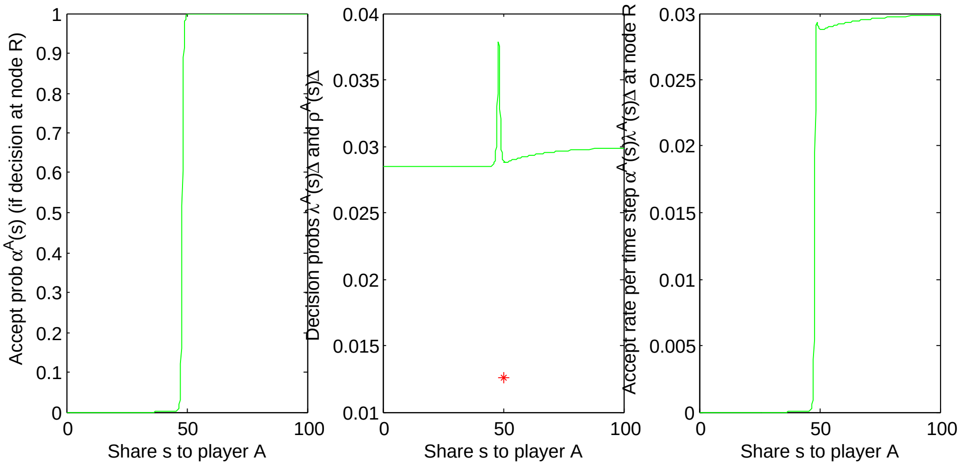

Note: Acceptance probabilities in rejection-first game, including errors in timing (choices described by Prop. 8). Green: Node R; red: node N.

Left panel: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

Middle panel: Decision arrival probabilities of player A: $\lambda^A(s)\Delta$ (green) and $\rho^A(s)\Delta$ (red).

Right panel: Acceptance probabilities $\alpha^A(s)\lambda^A(s)\Delta$ of player A.

compute bargaining equilibria again, under the the alternating and rejection-first protocols we considered previously, but we model the decisions according to the specification in Proposition 1 (Section 2.1), where there are no errors in timing.[22]

Equilibrium of the alternating offers game, under the decision model of Prop. 1, is illustrated in Figs. 14-15. Equilibrium under the rejection-first protocol, under the decision model of Prop. 1, is shown in Figs. 16-17. In the bottom-right panels of Figs. 14 and 16, we see that 100% of time available is allocated to making the decision at hand, since in the model of Prop. 1 time has no inherent value and monitoring whether or not the decision should be made now requires no time.

However, the overall conclusion about the importance of timing errors is very simple: results with and without timing errors are very similar. In both cases, the allocation ends up being a split near 50/50 with probability very near one. Under the benchmark parameterization, the arrival rate of the decision is also similar in both cases. Thus errors in timing have little effect, except that they make the model simpler to solve, since the possible corner solutions seen in Prop. 1 are eliminated when we move to the model analyzed in Prop. 8.[23] (Under an alternative

[22] A MATLAB program, onedecision_ns.m, is provided to solve the problem analyzed in Proposition 1. The program rubegame_alternating.m can then be used to solve each backwards induction step, in order to calculate the equilibria reported here. When calling rubegame_alternating.m, select DECISIONMODEL=1 to run the decision model without errors in timing (Prop. 1), or DECISIONMODEL=2 to run the decision model with errors in timing (Prop. 8). The same parameter settings are applicable to the program rubegame_rejectfirst.m that computes the rejection-first protocol.

[23] We should not extrapolate to conclude that timing errors are never important in this decision model. In the author's related work, Costain and Nakov (2014), timing errors are crucial for nonneutrality of monetary

parameterization where control costs were lower, $\kappa = 0.02$, it was found that decision arrival rates were substantially higher under the model of Prop. 1 than under the model of Prop. 8, because with lower control costs much less time was devoted to decision-making under the model of Prop. 8, whereas under the model of Prop. 1 all time is devoted to choice by assumption. But qualitatively the results were otherwise unchanged.)

3.4 Comparative statics

Figures 18-19 illustrate a comparative statics exercise in which we make time inherently more valuable, raising the productivity of alternative activities from its benchmark value of $A = 1$ to $A = 2$. (The specification is otherwise identical to the rejection-first benchmark of Figs. 12-13.) Since time is more costly, less of it is devoted to the decision. Comparing the red cross in the bottom-right panel of Fig. 18 to that of 18, the time devoted to alternative activities rises from only $h = 0.0.035$ when $A = 1$ to $h = 0.0397$ when $A = 2$. Therefore the decision arrives more slowly, and with less accuracy. In particular, the probability that A offers $s \in [49, 52]$ falls from 93.4% in the benchmark specification to 86.6% when $A = 2$. At the same time, the arrival rate of the first offer falls from 1.3% per time step in the benchmark specification to 0.9% per time step when $A = 2$ (red star in second panel of 19), and the arrival rate of a decision to reject an offer falls from roughly 3% per time step to roughly 1.6% per time step.

Thus, the main impact of making time in alternative activities more valuable is to slow down the arrival of the decision. Similar effects are observed in several other comparative statics exercises (these additional figures are not shown). In particular, with $A = 1$, lowering the noise parameter from $\kappa = 0.2$ to $\kappa = 0.02$ increases the arrival rate of initial offers from 1.3% per time step to 6.4% per time step and shrinks the range of offers made in equilibrium. Likewise, lowering the benchmark arrival rate $\bar{\rho}$ from 0.5 to 0.1 slows down the arrival rate of initial offers, from 1.3% per time step to 0.95% per time step, while raising $\bar{\rho}$ to 1 raises the arrival rate of initial offers to 1.5% per time step.

A rather different comparative statics exercise is considered in Figs. 20-21, in which the value of alternative activities is raised to $A = 3$. The importance of this example is that now the value of alternative activities is so high that the players would be better off never ending the game by splitting the "rotten pie". That is, $A/\delta = 60$ exceeds the value of half the pie (50) in this example.

The results are sensible. If these players were fully rational, they would simply avoid negotiating altogether. Being imperfectly rational, they do devote some time to the bargaining game, but we see from the red dot the bottom-right panel of Fig. 20 that almost all time is devoted to alternative activities ($h = 0.95$). Therefore offers arrive very slowly; the arrival rate of initial offers is 0.45% per time step (red star in Fig. 21). Since ending the game by accepting a 50/50 share is less valuable than continuing without agreement, the probability of offering accepting equal shares is very close to zero; A's acceptance probability jumps close to one only when an offer $s \geq 60$ is received. Since A knows that B is unlikely to accepts offers above $s = 40$, when A does make offers some of these are as generous as $s = 42.5$ (the smallest offer that is rejected almost certainly by B).

shocks, because they weaken the "selection effect" described by Golosov and Lucas (2007). But in the context of sequential bargaining, while time-consuming decisions are crucial for generating bargaining delays, the additional layer of errors on the timing itself has little impact on the equilibrium.

3.5 Multiplicity of equilibrium

The alternating offers game of Rubinstein (1982) implies a unique equilibrium with immediate agreement. In Perry and Reny's (1993) game with fixed time costs for decisions, agreement is not immediate, and multiple equilibria occur for two distinct reasons. First, multiplicity may arise due to the possibility of exactly simultaneous actions. Player i may make an offer at time t which she expects will be accepted by $-i$; this is supported as an equilibrium if $-i$ expects that his offer at t would be rejected. Since the same argument can be made for either player, which player offers at a given point in time is not uniquely determined. Second, the game of Perry and Reny displays multiple equilibria in which one of the agents expects, and on average receives, a higher share than the other. Thus, even under a perfectly symmetric parameterization, the expected payoffs of A and B need not be equal.

Multiplicity of the first type is ruled out almost by construction in our model, since precise control of timing is costly. While it is feasible in this model to make a proposal with probability one in a discrete time step Δ, such a rapid choice achieves minimal precision but requires a high fraction of time devoted to monitoring, the more so as Δ approaches zero. Normally, in equilibrium at a node where no offer is outstanding, both players are simultaneously thinking about possible offers. Knowing that i might complete an offer over the interval $[t, t+\Delta]$ with a probability that is positive but small (if Δ is small) does not suffice to rule out decision-making effort by player $-i$.

The issue of multiplicity of expected shares is separate from the possibility of multiplicity in timing. To investigate this question, we compute our benchmark specification from different initial guesses. On one hand, we compute it from a symmetric initial guess, $N^A_{init} = N^B_{init} = 50$, which converges to the symmetric equilibrium shown in Sec. 3.2, where $N^A_{ss} = N^B_{ss} = 48.181$. On the other hand, we start with a highly asymmetric initial guess, $N^A_{init} = 80$, $N^B_{init} = 20$. In this case, the asymmetry in players' expected shares unravels as backwards induction proceeds, converging finally to precisely the same equilibrium, with $N^A_{ss} = N^B_{ss} = 48.181$. Figure 22 plots the value functions associated with offers made and offers received, as calculated both simulations. The simulation that starts from a symmetric guess is plotted with M^A_{ss} in blue and R^A_{ss} in green. The simulation that starts from an asymmetric guess is plotted both functions in red, but the red lines are invisible because they are exactly overlaid by the blue and green lines.

Unravelling from the asymmetric initial guess happens very slowly. When player A expects to receive a very high share (represented by a high value of $N^A_{t+\Delta}$), that player will not accept low offers, so N^A_t cannot be much below the assumed value of $N^A_{t+\Delta}$. Indeed, if offers are constrained to a sufficiently coarse discrete grid, backwards induction may get "trapped" at a distribution of offers that remains asymmetrically favorable to one of the players. In other words, on a coarse grid, multiple equilibria may be sustained. This is illustrated in the left panel of Figure 23, which shows the equilibrium value functions M^A_{ss} and R^A_{ss} that are calculated from an initial guess that is extremely favorable to A, conditional on two different discrete grids. The dotted lines assume a grid of possible shares from 0 to 100 by steps of 2.5; in this case backwards induction converges to $N^A_{ss} = 53.1$, $N^B_{ss} = 43.6$. Given a step size of 1.0, the computation converges to $N^A_{ss} = 50.0$, $N^B_{ss} = 46.4$, but this is not shown, in order to avoid cluttering the graph. Finally, the solid line shows the benchmark equilibrium, which is computed with a step size of 0.1; this grid is sufficiently fine that backwards induction converges to a fully symmetric equilibrium with $N^A_{ss} = N^B_{ss} = 48.181$, even when we start from an asymmetric guess.[24]

[24]The solid curves are identical to those shown in Fig. 22 and in the top-left panel of Fig. 12.

The right panel of Fig. 23 shows a similar exercise under a higher degree of rationality, lowering the decision noise parameter from its benchmark value of $\kappa = 0.2$ to $\kappa = 0.02$. We see that when decisions are less costly, a wider range of equilibrium shares can be sustained, conditional on a given discrete grid. In this case, with a step size of 2.5, backwards induction from a guess very favorable to A converges to $N_{ss}^A = 76.7$, $N_{ss}^B = 22.4$ (dots). With step size 0.5, the same calculation converges to $N_{ss}^A = 62.1$, $N_{ss}^B = 37.3$ (dashes), and with step size 0.25, it converges to $N_{ss}^A = 54.8$, $N_{ss}^B = 44.6$ (dash-dot). Finally, with step size 0.1 it converges all the way to a symmetric equilibrium with $N_{ss}^A = N_{ss}^B = 49.66$ (shown as a solid line). Note that with a lower noise parameter ($\kappa = 0.02$), the initial value of the game is higher than it was under the benchmark parameterization ($\kappa = 0.2$), reflecting the fact that an agreement arrives faster and with less decision effort.

In all the cases shown in the two panels of the figure, a fully symmetric equilibrium is found if we start from a symmetric guess. But as the figure shows, on a coarse grid, the algorithm can also converge to an asymmetric equilibrium. As long as the grid is sufficiently fine, backwards induction converges to a unique equilibrium, which is symmetric as long as the two players have symmetric parameters. It is reasonable to conjecture that equilibrium is unique if the game is defined on a continuous space of possible offers, but demonstrating this rigorously is beyond the scope of the present paper.

4 Conclusions

This paper has presented a model of decision-making in which choices are error-prone because achieving perfect precision is excessively costly. The cost of choice is the time it requires. The decision-maker may choose quickly, achieving low precision, or more slowly, achieving arbitrarily high precision. The problem is defined in terms of an arbitrary discrete time step, and is well-behaved both in discrete time and in the continuous-time limit. Likewise, although the choice is defined across a discrete set of possible actions, it is well-behaved in the limit with a continuum of possible actions.

Besides errors in choices across two or more feasible actions, the paper also extends the control cost equilibrium concept by studying errors in the timing of decisions. Several scenarios are considered, with costly choice across alternatives only, or costly choice of timing only, or both. In the main model, where both choices and the timing of those choices are costly, optimal decision-making is characterized by an interior solution that varies smoothly with changes in the parameters of the problem. This occurs because errors in the timing of choice eliminate corner solutions that would otherwise complicate the analysis.

When this model of decisions is applied inside a game, the time devoted to choice must be reflected in the extensive form. As an application, this paper considers a sequential bargaining game. Two players may make offers to split a pie; making an offer is a time-consuming, error-prone decision. A player that has received an offer may accept the offer, ending the game, or reject it, returning the game to a state where no offer is outstanding; this binary choice is also a time-consuming, error-prone decision. An alternative protocol is considered, in which rejecting an offer requires the formulation of an alternative offer; our decision framework implies that this is a more costly decision, and therefore more time goes by in equilibrium under this alternative protocol before an offer is accepted.

Numerical simulations of these bargaining games reveal an equilibrium structure very similar to that of Rubinstein (1982), except that agreement is not achieved instantaneously. Under a symmetric parameterization, there is a very high probability of an approximately even split of the pie. When no offer is outstanding, both players simultaneously devote part of their time to formulating offers. Players typically propose shares very slightly favorable to themselves, and typically accept shares very slightly unfavorable to themselves. Sometimes players propose shares excessively favorable to themselves, which are typically rejected; proposing shares that are excessively generous to the opponent is a much more costly mistake and therefore occurs less frequently in equilibrium.

Perry and Reny (1993) also considered a sequential bargaining game with time-consuming formulation of offers. Their game exhibits two types of multiplicity of equilibrium. There may be multiplicity in the identity of the player that makes an offer at a given point of time; there may also be multiplicity in the expected equilibrium share received by a given player. In contrast, in the present paper, as long as the game is computed with a sufficiently short time step, and a sufficiently small step size between possible offers, equilibrium is unique. Assuming higher noise in decision-making also makes uniqueness more likely. Thus we find that the results of the Rubinstein (1982) and Binmore *et al.* (1983) games appear more robust than Perry and Reny's generalization suggested.

References

[1] Álvarez, Fernando; Francesco Lippi, and Luigi Paciello (2011), "Optimal price setting with observation and menu costs." *Quarterly Journal of Economics* 126 (4), pp. 1909-1960.

[2] Anderson, Simon; Jacob Goeree, and Charles Holt (2002), "The logit equilibrium: a perspective on intuitive behavioral anomalies." *Southern Economic Journal* 69 (1), pp. 21-47.

[3] Anderson, Simon; André de Palma, and Jacques-François Thisse (1992), *Discrete Choice Theory of Product Differentiation.* MIT Press.

[4] Ash, Robert (1965), *Information Theory.* Dover Publications.

[5] Baron, Richard; Jaques Durieu, Hans Haller, and Philippe Solal (2002), "Control costs and potential functions for spatial games." *International Journal of Game Theory* 31, pp. 541-561.

[6] Barro, Robert (1977), "Long-term contracting, sticky prices, and monetary policy." *Journal of Monetary Economics* 3, pp. 305-316.

[7] Binmore, Kenneth; Ariel Rubinstein, and Asher Wolinsky (1986), "The Nash bargaining solution in economic modelling." *Rand Journal of Economics* 17 (2), pp. 176-188.

[8] Calvo, Guillermo (1983), "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics* 12, pp. 383-98.

[9] Cheremukhin, Anton; Paulina Restrepo-Echevarria; and Antonella Tutino (2012), "The assignment of workers to jobs with endogenous information selection." Manuscript, Federal Reserve Bank of Dallas.

[10] Christiano, Lawrence; Martin Eichenbaum, and Charles Evans (2005), "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy* 113 (1), pp. 1-45.

[11] Costain, James, and Anton Nakov (2014), "Logit price dynamics." ECB Working Paper #1693.

[12] Cover, Thomas, and Joy Thomas (2006), *Elements of Information Theory*, 2nd ed. Wiley Interscience.

[13] Gertler, Mark; Luca Sala, and Antonella Trigari (2008), "An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining." *Journal of Money, Credit and Banking* 40 (8), 1713-1764.

[14] Goeree, Jacob, and Charles Holt (1999), "Stochastic game theory: for playing games, not just for doing theory." *Proc. Nat. Acad. Sci. USA* 96, pp. 10564-7.

[15] Goeree, Jacob, and Charles Holt (2001), "Ten little treasures of game theory and ten intuitive contradictions." *American Economic Review* 91 (5), 1402-1422.

[16] Hall, Robert (2005), "Employment fluctuations with equilibrium wage stickiness." *American Economic Review* 95 (1), pp. 50-65.

[17] Hall, Robert, and Paul Milgrom (2008), "The limited influence of unemployment on the wage bargain." *American Economic Review* 98 (4), pp. 1653-1674.

[18] Robert D. Luce (1959), *Individual Choice Behavior: a Theoretical Analysis.* Wiley, New York.

[19] Machina, Mark (1985), "Stochastic choice functions generated from deterministic preferences over lotteries.' ' *Economic Journal* 95, pp. 575-594.

[20] Magnani, Jacopo, Aspen Gorry, and Ryan Oprea (2013), "Time and state dependence in an Ss decision experiment." Manuscript, Univ. of British Columbia.

[21] Mankiw, N. Gregory (1985), "Small menu costs and large business cycles: a macroeconomic model of monopoly." *Quarterly Journal of Economics* 100 (2), pp. 529-37.

[22] Matejka, Filip, and Alisdair McKay (2011), "Rational inattention to discrete choices: a new foundation for the multinomial logit model." CERGE-EI Working Paper Series 442.

[23] Matejka, Filip (2011), "Rationally inattentive seller: sales and discrete pricing." Manuscript, CERGE-EI.

[24] Matejka, Filip, and Alisdair McKay (2013), "Simple market equilibria with rationally inattentive consumers." *American Economic Review* 2012, 102(3), pp. 24-29.

[25] Mattsson, Lars-Goram, and Jorgen Weibull (2002), "Probabilistic choice and procedurally bounded rationality." *Games and Economic Behavior* 41, pp. 61-78.

[26] McKelvey, Richard, and Thomas Palfrey (1995), "Quantal response equilibrium for normal form games." *Games and Economic Behavior* 10, pp. 6-38.

[27] McKelvey, Richard, and Thomas Palfrey (1998), "Quantal response equilibrium for extensive form games." *Experimental Economics* 1, pp. 9-41.

[28] Merlo, Antonio, and Charles Wilson (1995), "A stochastic model of sequential bargaining with complete information." *Econometrica* 63 (2), pp. 371-399.

[29] Merlo, Antonio, and Charles Wilson (1998), "Efficient delays in a stochastic model of bargaining." *Economic Theory* 11 (1), pp. 39-55.

[30] Merlo, Antonio, and Xun Tang (2010), "Identification of stochastic sequential bargaining models." Manuscript, Univ. of Pennsylvania.

[31] Perry, Motty, and Philip Reny (1993), "A non-cooperative bargaining model with strategically timed offers." *Journal of Economic Theory* 59, pp. 50-77.

[32] Ramsey, Kristopher, and Curtis Signorino (2009), "A statistical model of the ultimatum game." Manuscript, Princeton Univ.

[33] Rubinstein, Ariel (1982), "Perfect equilibrium in a bargaining model." *Econometrica* 50, pp. 97-109.

[34] Rubinstein, Ariel (2013), "Response time and decision-making: an experimental study." *Judgment and Decision Making* 8 (5), pp. 540-551.

[35] Selten, Reinhard (1975), "A reexamination of the perfectness concept for equilibrium points in extensive games." *International Journal of Game Theory* 4, pp. 25-55.

[36] Shannon, Claude (1948), "A mathematical theory of communication." *The Bell System Technical Journal* 27, pp. 379-423 and 623-56.

[37] Shimer, Robert (2005), "The cyclical behavior of equilibrium unemployment and vacancies." *American Economic Review* 95 (1), pp. 25-49.

[38] Shimer, Robert (2007), "Mismatch." *American Economic Review* 97 (4), pp. 1074-1101.

[39] Shimer, Robert, and Lones Smith (2000), "Assortative matching and search." *Econometrica* 68 (2), pp. 343-369.

[40] Sims, Christopher (1998), "Stickiness." *Carnegie-Rochester Conference Series on Public Policy* 49, pp. 317-56.

[41] Sims, Christopher (2003), "Implications of rational inattention." *Journal of Monetary Economics* 50, pp. 665-90.

[42] Smets, Frank, and Rafael Wouters (2003), "An estimated stochastic dynamic general equilibrium model of the euro area." *Journal of the European Economic Association* 1 (5), pp. 1123-1175.

[43] Stahl, Dale (1990), "Entropy control costs and entropic equilibrium. " *International Journal of Game Theory* 19, pp. 129-138.

[44] Theodorou, Evangelos; Krishnamurthy Dvijotham, and Emanuel Todorov (2013), "From information-theoretic dualities to path integral and Kullback Leibler control: continuous and discrete time formulations." Manuscript presented at 16th Yale Workshop on Learning and Adaptive Systems.

[45] Todorov, Emanuel (2009), "Efficient computation of optimal actions." *Proceedings of the National Academy of Sciences* 106 (28), pp. 11478-11483.

[46] Van Damme, Eric (1991), *Stability and Perfection of Nash Equilibrium*, 2nd edition. Springer Verlag.

[47] Wolinsky, Asher (1987), "Matching, search, and bargaining." *Journal of Economic Theory* 42, pp. 311-333.

[48] Woodford, Michael (2008), "Information-constrained state-dependent pricing." *Journal of Monetary Economics* 56, pp. S100-S124.

[49] Woodford, Michael (2013), "Perceptual coding and stochastic choice." Manuscript, Columbia University.

5 Appendix. Additional figures.

5.1 Games without errors in timing.

Figure 14: Equilibrium of alternating offers game, with no errors in timing.

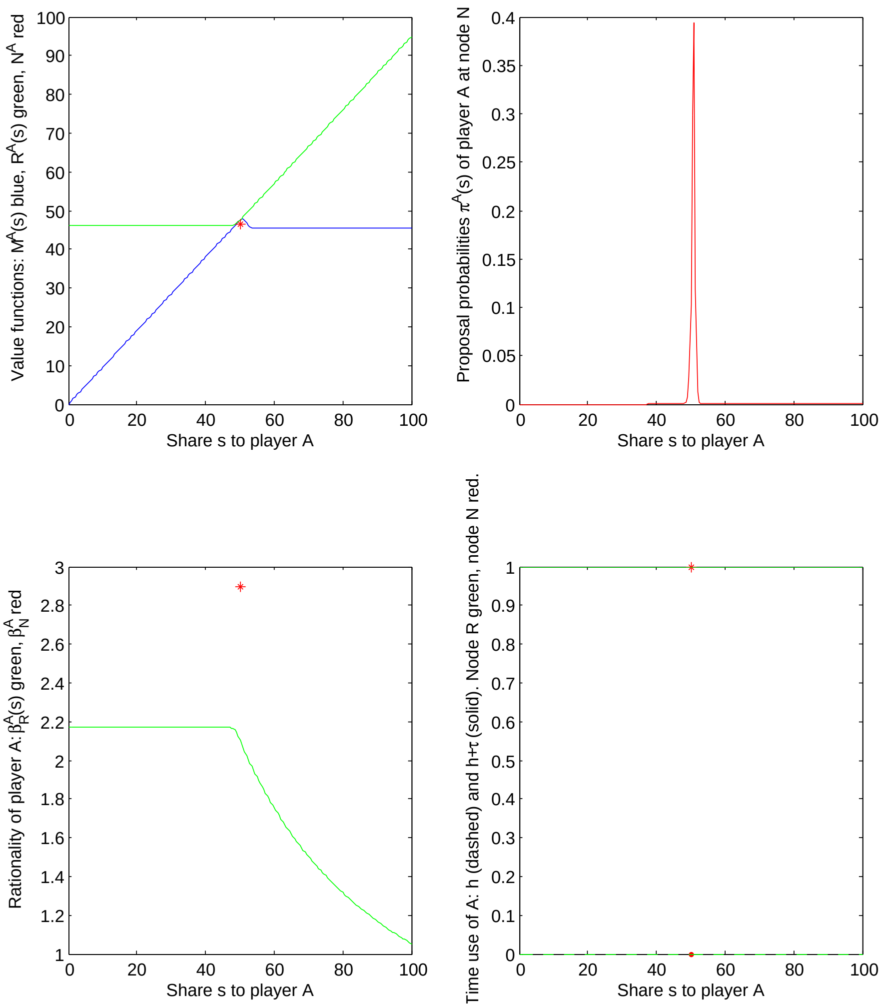

Note: Equilibrium of alternating offers game, with no errors in timing (choices described by Prop. 1). Green: Node R; blue: node M; red: node N.

Top left: Values to A of outstanding offers. Blue: Value $M^A(s)$ of offer made by A. Green: Value $R^A(s)$ of offer made by B. Red: Value N^A when no offer has been made.
Top right: Offer probabilities $\pi^A(s)$ when A makes an offer at node N.
Bottom left: Rationality of A. Green: rationality $\beta^A_R(s)$ when responding to offer s. Red: Rationality β^A_N when making an offer at node N.
Bottom right: Time use of A at nodes $R^A(s)$ and N^A. Plus: labor h; star: $h+\tau$.

Figure 15: Acceptance decision in alternating offers game, with no errors in timing.

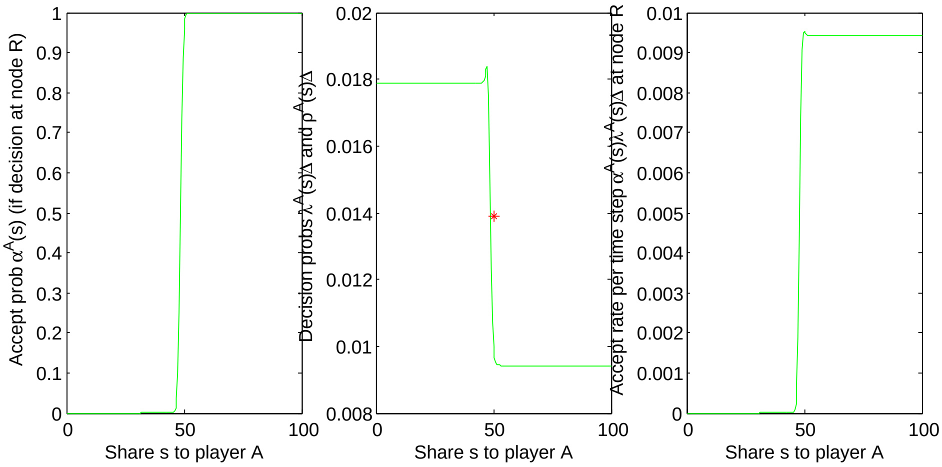

Note: Acceptance probabilities in alternating offers game, with no errors in timing (choices described by Prop. 1). Green: Node R; red: node N.

Left panel: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

Middle panel: Decision arrival probabilities of player A: $\lambda^A(s)\Delta$ (green) and $\rho^A(s)\Delta$ (red).

Right panel: Acceptance probabilities $\alpha^A(s)\lambda^A(s)\Delta$ of player A.

Figure 16: Equilibrium of rejection-first game, with no errors in timing.

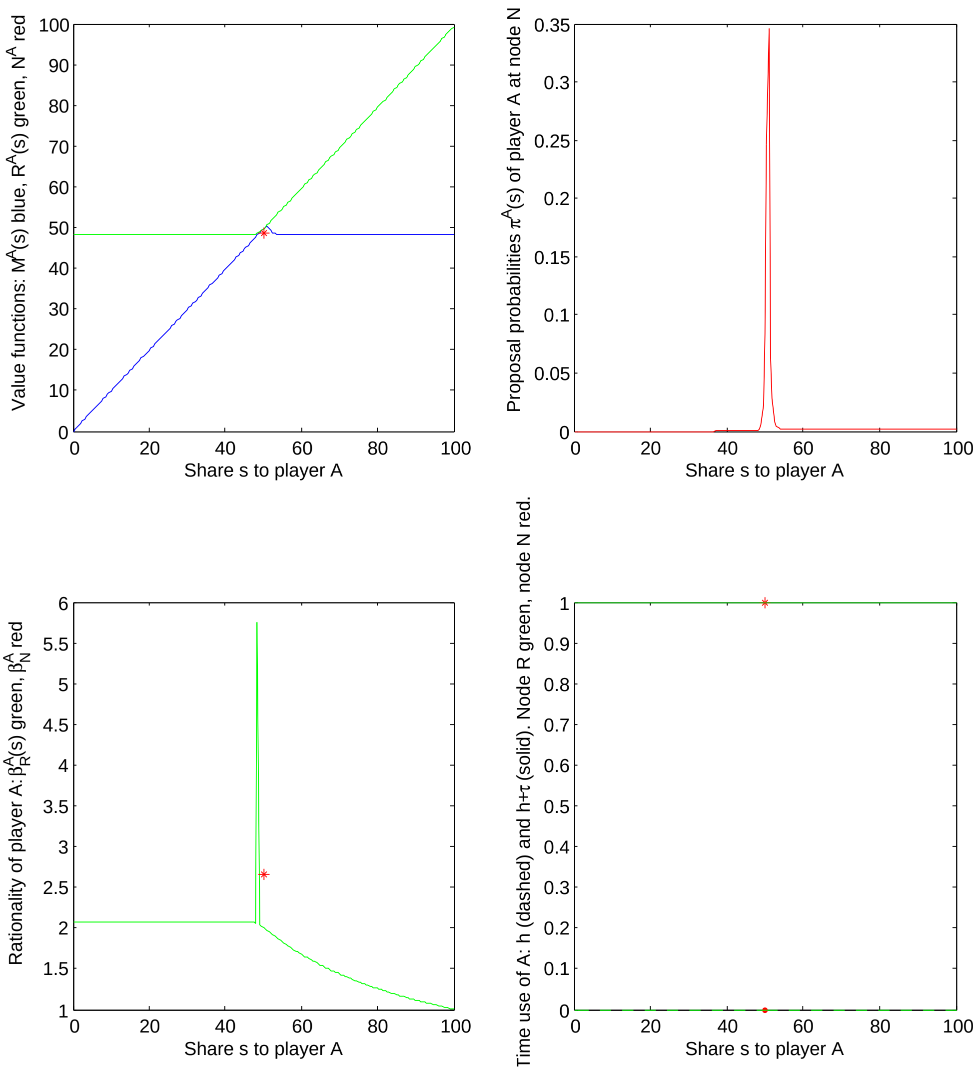

Note: Equilibrium of rejection-first game, with no errors in timing (choices described by Prop. 1). Green: Node R; blue: node M; red: node N.

Top left: Values to A of outstanding offers. Blue: Value $M^A(s)$ of offer made by A. Green: Value $R^A(s)$ of offer made by B. Red: Value N^A when no offer has been made.
Top right: Offer probabilities $\pi^A(s)$ when A makes an offer at node N.
Bottom left: Rationality of A. Green: rationality $\beta^A_R(s)$ when responding to offer s. Red: Rationality β^A_N when making an offer at node N.
Bottom right: Time use of A at nodes $R^A(s)$ and N^A. Plus: labor h; star: $h + \tau$.

Figure 17: Acceptance decision in rejection-first game, with no errors in timing.

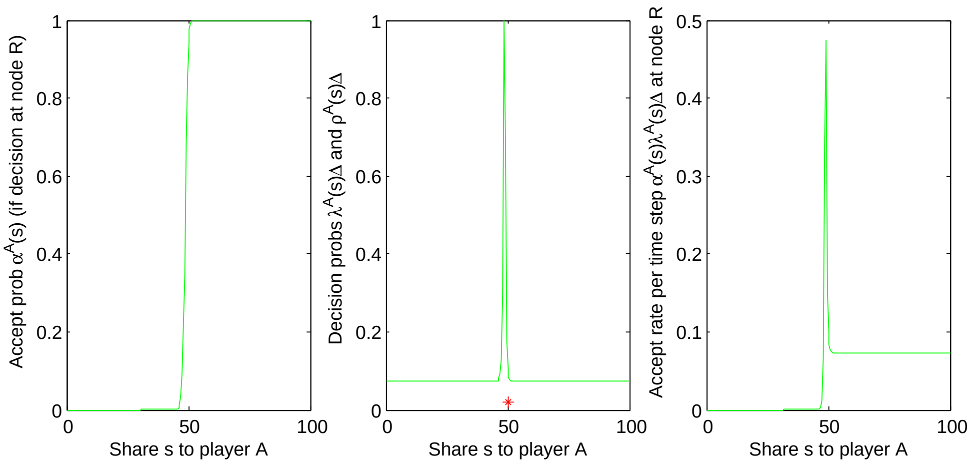

Note: Acceptance probabilities in rejection-first game, with no errors in timing (choices described by Prop. 1). Green: Node R; red: node N.

Left panel: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

Middle panel: Decision arrival probabilities of player A: $\lambda^A(s)\Delta$ (green) and $\rho^A(s)\Delta$ (red).

Right panel: Acceptance probabilities $\alpha^A(s)\lambda^A(s)\Delta$ of player A.

5.2 Comparative statics

Figure 18: Equilibrium of rejection-first game, assuming costlier time ($A = 2$).

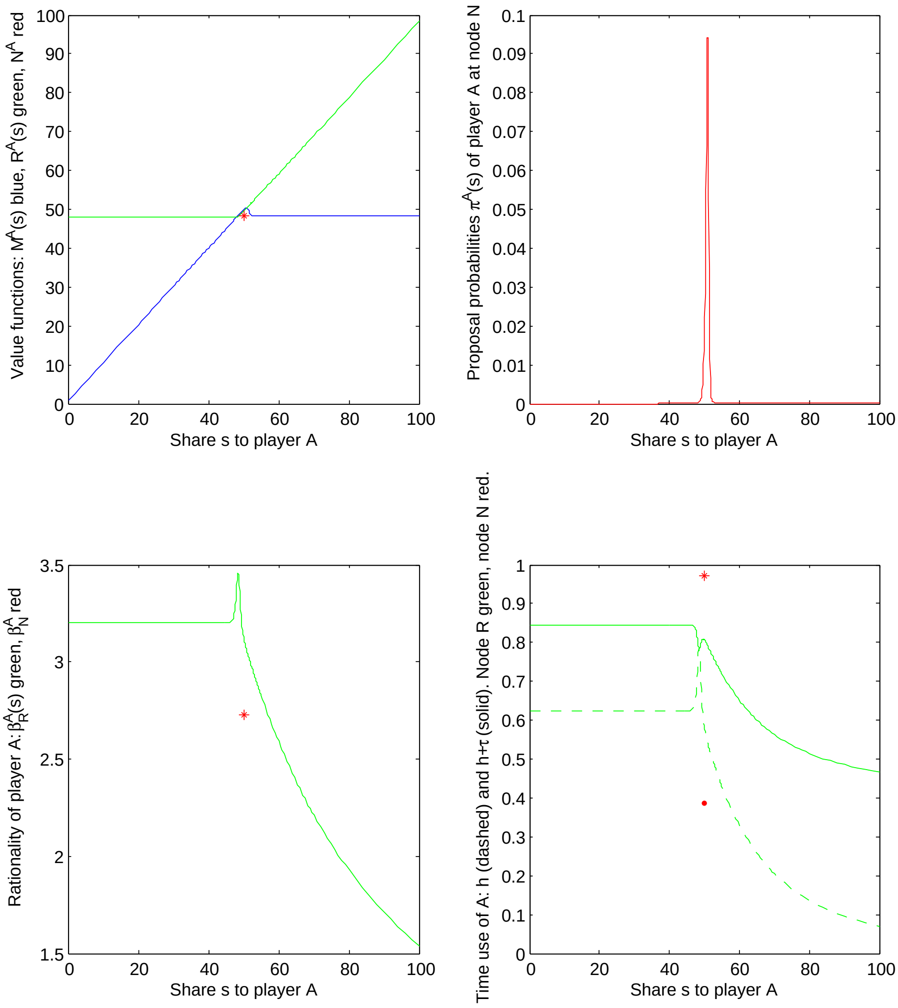

Note: Equilibrium of rejection-first game, including errors in timing (choices described by Prop. 8). Green: Node R; blue: node M; red: node N. Assuming costlier time ($A = 2$ instead of $A = 1$).

Top left: Values to A of outstanding offers. Blue: Value $M^A(s)$ of offer made by A. Green: Value $R^A(s)$ of offer made by B. Red: Value N^A when no offer has been made.
Top right: Offer probabilities $\pi^A(s)$ when A makes an offer at node N.
Bottom left: Rationality of A. Green: rationality $\beta_R^A(s)$ when responding to offer s. Red: Rationality β_N^A when making an offer at node N.
Bottom right: Time use of A at nodes $R^A(s)$ and N^A. Plus: labor h; star: $h + \tau$.

Figure 19: Acceptance decision in rejection-first game, assuming costlier time ($A = 2$).

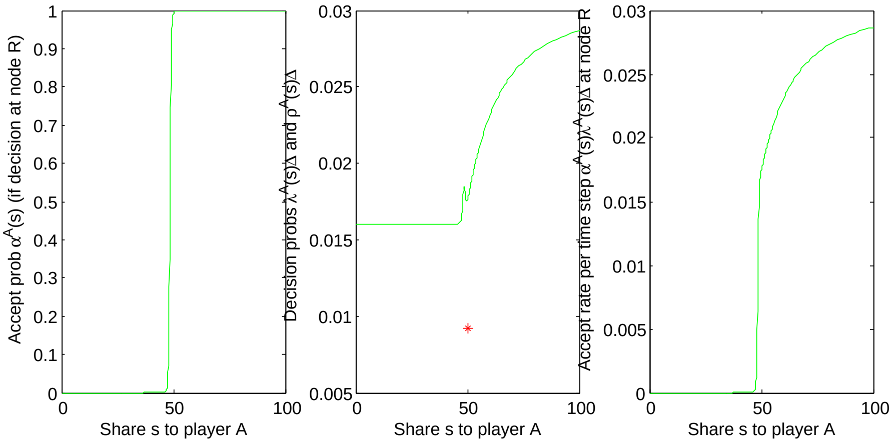

Note: Acceptance probabilities in rejection-first game, including errors in timing (choices described by Prop. 8). Green: Node R; red: node N. Assuming costlier time ($A = 2$ instead of $A = 1$).

Left panel: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.
Middle panel: Decision arrival probabilities of player A: $\lambda^A(s)\Delta$ (green) and $\rho^A(s)\Delta$ (red).
Right panel: Acceptance probabilities $\alpha^A(s)\lambda^A(s)\Delta$ of player A.

5.3 Sharing a rotten pie

Figure 20: Equilibrium of rejection-first game, assuming "rotten pie" ($A = 3$).

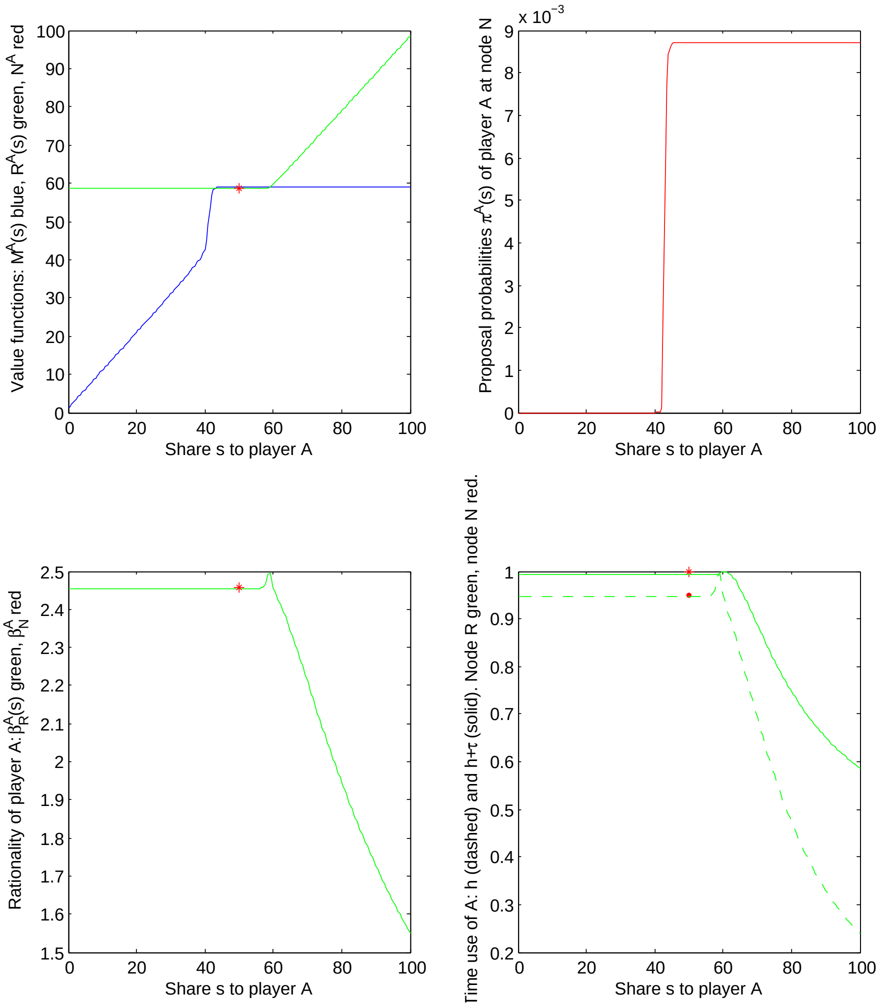

Note: Equilibrium of rejection-first game, including errors in timing (choices described by Prop. 8). Green: Node R; blue: node M; red: node N. Time value raised to $A = 3$, implying "rotten pie": $50 < A/\delta$.

Top left: Values to A of outstanding offers. Blue: Value $M^A(s)$ of offer made by A. Green: Value $R^A(s)$ of offer made by B. Red: Value N^A when no offer has been made.
Top right: Offer probabilities $\pi^A(s)$ when A makes an offer at node N.
Bottom left: Rationality of A. Green: rationality $\beta^A_R(s)$ when responding to offer s. Red: Rationality β^A_N when making an offer at node N.
Bottom right: Time use of A at nodes $R^A(s)$ and N^A. Plus: labor h; star: $h + \tau$.

Figure 21: Acceptance decision in rejection-first game, assuming “rotten pie” ($A = 3$).

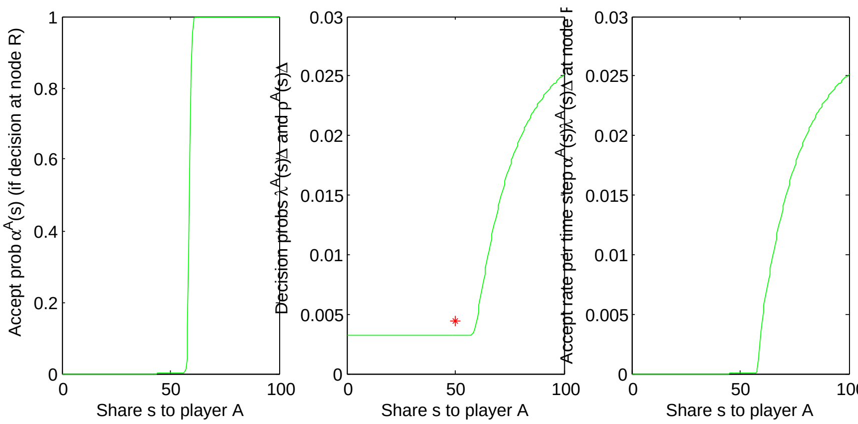

Note: Acceptance probabilities in rejection-first game, including errors in timing (choices described by Prop. 8). Green: Node R; red: node N. Time value raised from $A = 1$ to $A = 3$, implying “rotten pie”: $50 < A/\delta$.

Left panel: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.
Middle panel: Decision arrival probabilities of player A: $\lambda^A(s)\Delta$ (green) and $\rho^A(s)\Delta$ (red).
Right panel: Acceptance probabilities $\alpha^A(s)\lambda^A(s)\Delta$ of player A.

5.4 Multiplicity of equilibrium

Figure 22: Equilibrium of rejection-first game. Example of uniqueness.

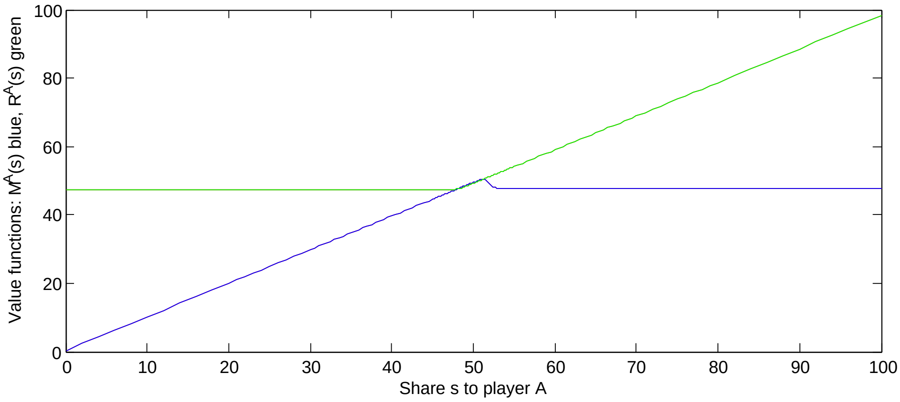

Note: Functions $R^A(s)$ and $M^A(s)$ of rejection-first game, including timing errors (choices described by Prop. 8). Green: Node R; blue: node M. Parameters identical to benchmark equilibrium shown in Fig. 12.

Note: Two equilibrium simulations are shown, assuming step size 0.1 in offer space.
Equilibrium calculated from asymmetric starting guess is plotted in red.
Equilibrium calculated from symmetric starting guess is plotted in green and blue, but is numerically identical and therefore covers up the red lines.

Figure 23: Equilibrium of rejection-first game. Examples of multiplicity.

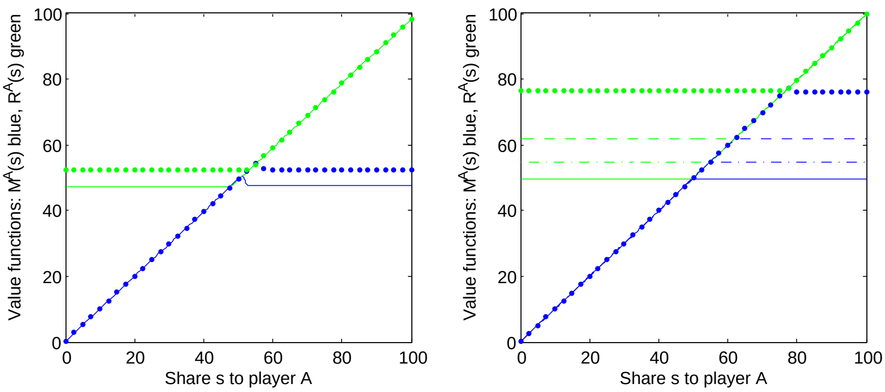

Note: Functions $R^A(s)$ and $M^A(s)$ of rejection-first game, including timing errors (choices described by Prop. 8). Green: Node R; blue: node M. Parameters identical to benchmark equilibrium, except for κ and step size in offer grid. Initial guess is very favorable to A ($N^A_{init} >> N^B_{init}$).

Left panel: High noise: $\kappa = 0.2$. Dots: step size 2.5. Solid: step size 0.1.
Right panel: Low noise: $\kappa = 0.02$. Dots: step 2.5. Dashes: step 0.5. Dot-dash: step 0.25. Solid: step 0.1.
In both panels, equilibrium converges to symmetry when step size is 0.1, but remains asymmetric in the other examples.