Costly decisions and sequential bargaining

James Costain
Banco de España

July 2017

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, thus endogenizing the order and timing of offers and responses in the game.

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 and offers are sometimes rejected. In contrast to the model of Perry and Reny, our numerical results indicate that equilibrium is unique when the space of possible offers is sufficiently finely spaced.

Keywords: Bargaining, control costs, logit equilibrium, near-rational choice.
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]The author is grateful for helpful discussions with Henrique Basso, Filip Matejka, Pascal Michaillat, Espen Moen, Anton Nakov, Plamen Nenov, Galo Nuño, Ariel Rubinstein, Antonella Trigari, and Ernesto Villanueva, and from seminar participants at the Banco de España, the Univ. of Murcia, UC Santa Cruz, Univ. Carlos III, CEF 2014, SaM 2015, ESSIM 2015, CEF 2015, and EEA 2015. 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 (Moreno and Wooders 1998; 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. A decision is conceived as a random variable distributed over a set of possible actions, and the cost of the decision is assumed to increase with the precision of this random variable. The player maximizes the payoffs that would obtain in a costlessly rational game, net of the decision costs. Usually, a player prefers to spread 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 sophisticated enough to consider the costs and limitations of their own rationality when they make choices — an appealing property.

Realistically, *time use* is likely to be an important component of the costs of choice. Time devoted to decision-making may have an opportunity cost, and also implies discounting of the terminal payoffs received when the choice is made. In the context of the independent choices of a single decision-maker, which resources are consumed by the choice process may be irrelevant, since time, money, and other resources may be fungible. But time-consuming choice could be an important feature of a game, as it could have strategic implications, representing an opportunity for other players to intervene with actions of their own. Therefore, the time used up by decisions *should be reflected in the extensive form of the game.* Accordingly, this paper explores the implications of control cost equilibrium when we take the role of time in decision-making seriously.

The model, developed in Section 2, studies a decision-maker (DM) who may 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 then takes the form of a multinomial logit (as in Mattsson and Weibull, 2002).

Note, however, that in many dynamic situations there may exist a default option that occurs in the absence of any deliberate action. Our model allows for situations of this type by treating the *timing* of any deliberate choice as a time-consuming, error-prone decision too. Under the functional forms assumed in the paper, decision timing 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 (Secs. 2.2-2.3). 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

[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, de Palma, and Thisse (1992).

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 embeds the model of costly decisions into a game where two 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, these decisions are distinct in our setup; Sections 3.1-3.3 compare several different bargaining protocols.

Sections 3.4-3.5 further explore bargaining equilibrium by numerical simulation. The equilibrium resembles that of the Binmore, Rubinstein, and Wolinsky (1986) framework, except that the time to reach agreement is nonzero, and offers are not always accepted. The simulations show that equilibrium changes in intuitive ways as underlying parameters vary; this is true even in an example where parameters imply that splitting the pie is a bad idea. 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 a sufficiently fine grid. 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 (1986). Wolinsky (1987) and Hall and Milgrom (2008) 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. Bono and Wolpert (2010) study a bargaining game in which offers are stochastic, either due to errors or due to other unmodeled variation across players. Merlo and Wilson (1995, 1998) and Merlo and Wilson (2010) study bargaining games where the pie evolves randomly over time.

This paper is also motivated by the author's previous work on price stickiness. Costain and Nakov (2015, henceforth CN15) 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. The present paper extends their framework by allowing for a nonlinear cost of time and shows how it can be used to endogenize the order and timing of actions in an extensive-form game.[4]

Another motivation for this paper is its potential for unifying the analysis of several different margins in labor markets. 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

[4]CN15 studied a model with a linear cost of time use, which made it possible to separate the decision process into two Bellman steps, one governing the decision whether to adjust the price, and another governing the choice of which new price to set conditional on adjustment. The simultaneous treatment of both decision margins discussed in Sec. 2.3 makes the model of the present paper compatible with a nonlinear cost of time.

are themselves, likewise, engaged in costly choice). The present paper assumes a similar cost function to the choices involved in a bargaining game, which could be applied to 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 may offer a single microfoundation for many 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 could provide empirical implications about how errors vary 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 McKay 2015). Similar cost functions have been used to model other limitations on rationality, as in the model uncertainty framework 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, Durieu, Haller, and Solal 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 Woodford (2008) and Cheremukhin, Restrepo-Echevarria, and Tutino (2012) papers 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, rational inattention and control costs 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 stages are likely to be costly, and both have the same primary implication: an imperfect correlation between a player's true state and her 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 is typically 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 Models of time-consuming decisions

2.1 One costly choice

We first analyze a single costly decision, which will subsequently be a building block for more complex games. We regard a decision as a random variable distributed across a list of possible alternatives i. The cost of making a choice is that it requires time, for thinking, calculating, or otherwise comparing options on the basis of available information.[5] The decision-maker faces a tradeoff: if she chooses more quickly, her decision will be less precise.[6] Precision is measured in terms of relative entropy (also known as Kullback-Leibler divergence), which is a measure of distance from some default probability distribution $\vec{\eta}$. When the choice probabilities $\vec{\pi}$ are set equal to the default distribution, the decision cost is zero; the time cost of any decision that deviates from the default probabilities is proportional to the Kullback-Leibler divergence. That is, the time cost of decision $\vec{\pi}$ is given by $\kappa\mathcal{D}(\vec{\pi}||\vec{\eta})$, where $\vec{\eta}$ is an exogenously-given benchmark distribution over the same alternatives.[7]

It would be inelegant, unrealistic, and inconvenient to assume that the time required to solve a decision problem is exactly known. So rather than assuming that the DM chooses the solution time directly, we assume that she controls the arrival rate, ρ, of the solution. The expected duration of the choice process is (approximately) $1/\rho$; a slower solution rate implies more time spent on the choice, and we assume this permits a more precise allocation of probability across the alternatives under consideration. Fig. 1 graphs the game tree, showing one discrete time step that begins at time t (time subscripts are suppressed). Although we describe the model in discrete time, for most applications the time step should be considered very brief.[8] The DM chooses the solution rate ρ_t at node C ("calculating"); the curve under node C indicates that the solution rate is chosen from the unit interval, $\rho_t \in [0, 1]$. Since the actual time to reach a conclusion is stochastic, we see that the solution arrives with probability ρ_t (the DM reaches the "deciding" node D), while the DM remains undecided with probability $1 - \rho_t$, in which case the game returns to a node of type C at time $t + 1$.

When Nature allows a conclusion to arrive, the DM allocates probabilities $\pi_{i,t}$ across the alternatives $i \in \Gamma \equiv \{1, 2, ...n\}$. The precision of the decision $\vec{\pi}_t \equiv (\pi_{1,t}, \pi_{2_t}, \ldots, \pi_{n,t})' \in \Delta^{n-1}$ is lower when the solution rate is high (since $\vec{\pi}_t$ is a probability vector of length n, it is chosen from the simplex of order $n - 1$, written as Δ^{n-1}). Concretely, we assume

$$\frac{1}{\rho_t} \;\geq\; \kappa\mathcal{D}(\vec{\pi}_t||\vec{\eta}) \;\equiv\; \kappa\left(\sum_{i=1}^{n} \pi_{i,t} \ln\left(\frac{\pi_{i,t}}{\eta_i}\right)\right). \tag{1}$$

That is, a slower decision can have greater precision, where precision is measured by the

[5] Note that under this interpretation, the DM is assumed to have sufficient information to evaluate the alternatives considered but is not assumed to know which option is best when the decision process begins. The presence of the values V_i of the alternatives i in Bellman equation (2), should not be taken to mean that the DM "knows" the values V_i. Instead, it can be interpreted as saying that the DM has adequate information to calculate each V_i, given sufficient time.

[6] The model of this subsection is closely related to the "errors-in-pricing" specification considered in CN15, and to the models of Mattsson and Weibull (2002) and Matejka and McKay (2015): it allows for errors in a choice across a set of options, but ignores the possibility of any errors in timing.

[7] 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.

[8] The model has a well-behaved continuous-time limit, but this paper will focus on a discrete-time formulation.

Figure 1: A decision that requires time.

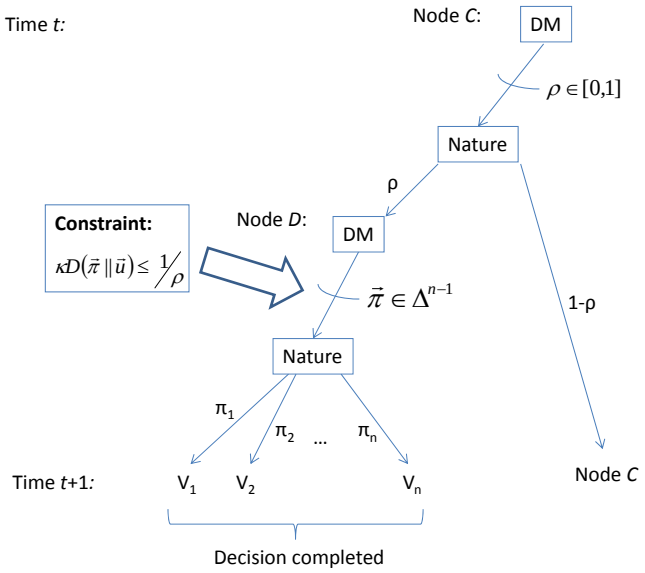

Note: Decision maker (DM) chooses solution rate ρ_t of decision; Nature lets conclusion arrive with probability ρ_t. A slower decision allows more precise allocation of probabilities $\pi_{i,t}$ across alternatives i with values $V_{i,t+1}$, satisfying the constraint $\kappa \mathcal{D}(\vec{\pi}_t || \vec{\eta}) \leq \rho_t^{-1}$. Time subscripts suppressed in diagram.

Kullback-Leibler divergence between the decision $\vec{\pi}_t$ and the benchmark distribution $\vec{\eta}$. Given the probabilities resulting from the DM's decision efforts, the actual option selected is random (chosen by nature); if option i is selected, the DM receives value $V_{i,t+1}$ at time $t+1$.

The model is written with discreteness in two dimensions: a choice set of n discrete options, and discrete time steps normalized to length one. 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 assumption is fundamentally just a matter of notational convenience.

The time required for more precise decisions could have a variety of costs. As in Rubinstein (1982), we first assume that the only cost is the pure time discounting, at rate δ, associated with delay. Now, let W_t denote the value of the problem at node C at time t. The value must

satisfy the following Bellman equation:

$$W_t = \max_{\rho, \{\pi_i\}_{i=1}^n} (1-\delta)\left[\rho \sum_{i=1}^n \pi_i V_{i,t+1} + (1-\rho) W_{t+1}\right] \tag{2}$$

$$\text{subject to:} \quad \rho\kappa \sum_{i=1}^n \pi_i \ln\left(\frac{\pi_i}{\eta_i}\right) \leq 1 \tag{3}$$

$$\text{and} \quad \sum_{i=1}^n \pi_i = 1. \tag{4}$$

Writing the multipliers on the two constraints as λ_ρ and λ_π, respectively, the first-order conditions are:

$$\rho(1-\delta)V_{i,t+1} - \lambda_\rho \rho\kappa\left(1 + \ln\left(\frac{\pi_i}{\eta_i}\right)\right) - \lambda_\pi = 0 \tag{5}$$

$$(1-\delta)\left(\sum_i \pi_i V_{i,t+1} - W_{t+1}\right) - \lambda_\rho \kappa \mathcal{D}(\vec{\pi}||\vec{\eta}) = 0 \tag{6}$$

$$1 - \rho\kappa\mathcal{D}(\vec{\pi}||\vec{\eta}) = 0 \tag{7}$$

$$1 - \sum_i \pi_i = 0. \tag{8}$$

Rearranging, the necessary condition for π_i can be rewritten as

$$\frac{\pi_i}{\eta_i} = \exp\left(\beta V_{i,t+1} - \frac{\beta\lambda_\pi}{\rho(1-\delta)} - 1\right), \tag{9}$$

where

$$\beta = \frac{(1-\delta)}{\kappa\lambda_\rho}. \tag{10}$$

Hence, π_i is proportional to the exponential of the value of alternative i, scaled by a coefficient β which represents the precision of the choice. Since the probabilities must sum to one, (9) is solved by a multinomial logit across the various alternatives i:

$$\pi_{i,t} = \frac{\eta_i \exp(\beta_t V_{i,t+1})}{\sum_j \eta_j \exp(\beta_t V_{j,t+1})}. \tag{11}$$

It is helpful now to introduce the notation $E^{\pi_t} x_{t+1} \equiv \sum_{i=1}^n \pi_{i,t} x_{i,t+1}$ to represent the time t expectation of a random variable x_{t+1} with distribution $\vec{\pi}_t$. Hence if $\vec{\eta}$ is the default distribution, we will write $E^\eta x = \sum_{i=1}^n \eta_i x_i$. Note that the optimal distribution (11) implies

$$\pi_{i,t}\ln\pi_{i,t} = (\beta_t V_{i,t+1} + \ln\eta_i)\,\pi_{i,t} - \pi_{i,t}\ln\left(E^\eta \exp(\beta_t V_{t+1})\right) \tag{12}$$

and therefore

$$\mathcal{D}(\vec{\pi}_t||\vec{\eta}) = \sum_{i=1}^n \pi_{i,t}\ln\pi_{i,t} - \sum_{i=1}^n \pi_{i,t}\ln\eta_i = \beta_t E^{\pi_t} V_{t+1} - \ln\left(E^\eta \exp(\beta_t V_{t+1})\right) \equiv K(\beta_t). \tag{13}$$

Here $E^{\pi_t} V_{t+1}$ is the expected value of finishing the decision, which is

$$E^{\pi_t} V_{t+1} \;=\; \sum_i \pi_{i,t} V_{i,t+1} \;=\; \frac{E^\eta V_{t+1} \exp(\beta_t V_{t+1})}{E^\eta \exp(\beta_t V_{t+1})}. \tag{14}$$

The function $K(\beta)$ represents the entropy cost of achieving precision level β when probabilities are allocated optimally across the alternatives, according to (11).

Using (10), the first-order condition (6) that determines the solution rate ρ_t can be rewritten in terms of equilibrium precision β_t, eliminating the multiplier λ_ρ:

$$\beta_t \left(E^{\pi_t} V_{t+1} - W_{t+1}\right) \;=\; \mathcal{D}(\vec{\pi}_t || \vec{\eta}). \tag{15}$$

But now we have two different equations for the cost measure $\mathcal{D}$: (13) expresses the time cost when the probabilities $\pi_{i,t}$ are allocated optimally across alternatives i, while (15) rewrites (6) as an optimal tradeoff between the solution cost $\mathcal{D}$ and the benefit of arriving at the solution, which is $E^{\pi_t} V_{t+1} - W_{t+1}$. Combining (13) and (15), both $\mathcal{D}$ and $\beta_t E^{\pi_t} V_{t+1}$ cancel, leaving one simple equation to determine β_t, given the values $\vec{V}_{t+1}$ and W_{t+1} of the alternatives and of continuation:

$$\beta_t W_{t+1} \;=\; \ln(E^\eta \exp(\beta_t V_{t+1})), \tag{16}$$

or equivalently,[9]

$$E^\eta \exp(\beta_t (V_{t+1} - W_{t+1})) \;=\; 1. \tag{17}$$

The function $g(\beta, \vec{V}) \equiv \ln(E^\eta \exp(\beta V))$ that appears on the right-hand side of (16) is called the *cumulant generating function* of the random variable V under distribution η. (It is the log of the moment generating function.) To clarify the behavior of the equation (16) that determines β, it is helpful to state some properties of the functions g and K.

Lemma 1 *Let* $g(\beta, \vec{V}) \equiv \ln(E^\eta \exp(\beta V)) \equiv \ln\left(\sum_{i=1}^n \eta_i \exp(\beta V_i)\right)$. *We have:*

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

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

(c.) $\frac{\partial^2 g}{\partial \beta^2}(\beta, \vec{V}) = \frac{E^\eta V^2 \exp(\beta V)}{E^\eta \exp(\beta V)} - \left(\frac{E^\eta V \exp(\beta V)}{E^\eta \exp(\beta V)}\right)^2 = E^\pi V^2 - (E^\pi V)^2$. *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* $V_i = E^\eta V$ *for all* i.

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

(e.) *Let* $K(\beta)$ *be given by (13). Then* $K'(\beta) \geq 0$ *for all* $\beta > 0$ *with equality if and only if* $V_i = E^\eta V$ *for all* i. *Also,* $K'(0) = 0$, *and* $\lim_{\beta \to \infty} K'(\beta) = 0$.

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

Points (b) and (c) follow directly by differentiating and then simplifying using the logit probability formula (11). 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 (11). Point (d) follows from the fact that $\exp(x)$ is a convex function. Therefore $E^\eta \exp(\beta V) \geq \exp(\beta E^\eta V)$, with equality only if all values V_i are equal, so that $V_i = E^\eta V$ for all i. Therefore if there is any difference across the values V_i, we have $\ln(E^\eta \exp(\beta V)) > \ln(\exp(\beta E^\eta V)) = \beta E^\eta V$.

[9] Equation (16), illustrated in Fig. 2, is analytically equivalent to (17). But (17) is easier to solve numerically if $W_{t+1} > 0$, because subtracting W_{t+1} from $V_{i,t+1}$ before exponentiating helps avoid numerical overflow.

Figure 2: Optimal choice of precision β_t under backwards induction.

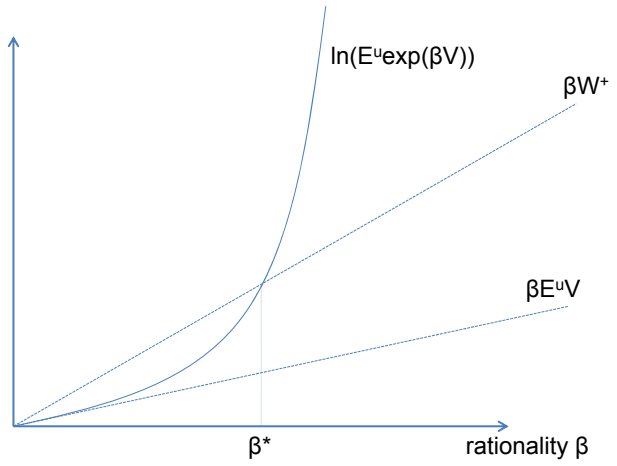

Note: The figure illustrates the interior solution stated in Prop. 1(b). Curve $g(\beta, \vec{V}_{t+1}) = \ln(E^\eta \exp(\beta V_{t+1}))$ is tangent to $\beta E^\eta V_{t+1}$ at $\beta = 0$, and has limiting slope $V^* \equiv \max_i V_{i,t+1}$ as $\beta \to \infty$ (the graph assumes $V^* > W_{t+1}$). If instead $V^* \leq W_{t+1}$, then $g(\beta, \vec{V}_{t+1})$ lies everywhere below βW_{t+1}, and postponement is optimal; see Prop. 1(a). If instead $E^\eta V_{t+1} \geq W_{t+1}$, the $g(\beta, \vec{V}_{t+1})$ lies everywhere above βW_{t+1}, and the corner $\beta_t = 0$ is optimal; see Prop. 1(c).

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 $K'(0) = 0$. Part (f) holds because $\pi_i \ln \pi_i \leq 0$ for all i, with equality if π_i only takes the value 0 or 1 for all i. Since K has a finite limit, we also note that $\lim_{\beta \to \infty} K'(\beta) = 0$. 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_t. We study the solution of problem (2) when W_{t+1} is taken as a given parameter, as is the case when we solve (2) by backwards induction. By Lemma 1(b), the curve $g(\beta, \vec{V}_{t+1})$, plotted as a function of β, has slope $E^\eta V_{t+1}$ at the origin. In the limit as $\beta \to \infty$, the slope of $g(\beta, \vec{V}_{t+1})$ converges to $V^* \equiv \max_i V_{i,t+1}$. Therefore, since g is convex, there exists exactly one positive β_t that solves (16) as long as $E^\eta V_{t+1} < W_{t+1} < V^*$. This β_t solves a given iteration step of (2), as long as the period length is sufficiently short.

Proposition 1 states this conclusion formally, and also describes the corner solutions that arise when the period length is longer, or when W_{t+1} lies outside of the bounds consistent with an interior solution.[10]

Proposition 1 *Consider problem (2)-(4), taking* W_{t+1} *and* $V_{i,t+1}$, $i \in \{1, \ldots, n\}$, *as given.*

[10]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 $\sum_i \exp(\beta^* V_{i,t+1}) > \infty^{NUM}$, where ∞^{NUM} is the largest real number representable on the computer. A similar caveat applies to the immediate solutions discussed in Prop. 1(c). When this arises in our simulated solutions, we set β^* to a large finite number that avoids overflow.

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

(b.) Interior solution: *Suppose* $V^* \equiv \max_i V_{i,t+1} > W_{t+1} > E^\eta V_{t+1}$, *and* $\kappa^{-1} < \ln n$. *Then there exists a unique positive* β^* *that satisfies (16). Let* $\mathcal{D}^*$ *and* ρ^* *be the entropy and arrival rate given by (13) and (18) when* $\beta = \beta^*$. *If* $\rho^* < 1$, *then (2)-(4) is solved by* $\beta_t = \beta^*$, $\rho_t = \rho^*$, *achieving the value*

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

(c.) Immediate solution: *Suppose* $E^\eta V_{t+1} \geq W_{t+1}$, *or that the conditions of part (b) hold, but the implied arrival probability in the first time step exceeds one:* $\rho^* \geq 1$. *Then the optimal arrival rate is* $\rho_t = 1$, *resulting in a precision level* β_t *that solves (19). Alternatively, suppose* $\kappa^{-1} \geq -\ln \eta^*$. *Then (2)-(4) is solved with infinite precision and arrival rate* $\rho_t = 1$.

The main point is that if $V^* > W_{t+1}$ and the time step is sufficiently short, then an interior solution applies, and precision is given by the unique positive β_t that solves (16). Knowing β_t, we can solve for the other endogenous quantities in the model: $\pi_{i,t}$ and $E^{\pi_t}V_{t+1}$ are given by (11) and (14), the entropy cost $\mathcal{D}_t$ is given by (13), and the arrival rate of the solution is

$$\rho_t \;=\; \frac{1}{\kappa \mathcal{D}(\vec{\pi}_t || \vec{\eta})}\,. \tag{18}$$

However, if the value V^* of the best possible option is less than the continuation value W_{t+1}, then the decision should simply be postponed. If instead $V^* > W_{t+1}$ and a precision level greater than or equal to the β^* that solves (16) can be achieved within a single time step, it is optimal to do so. Under this corner solution, the equilibrium precision β_t can be backed out from the time constraint. If β_t is finite, it is the unique solution to

$$\beta \frac{E^\eta V_{t+1} \exp(\beta V_{t+1})}{E^\eta \exp(\beta V_{t+1})} - \ln(E^\eta \exp(\beta V_{t+1})) \;=\; \kappa^{-1}. \tag{19}$$

The solution probability is then $\rho_t = 1$ in the first time step.[11]

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_{t+1}. Fig. 2 shows that given W_{t+1}, the solution β_t is *unaffected* by the information cost parameter κ. Likewise, given W_{t+1}, the probabilities $\vec{\pi}_t$, as well as $E^{\pi_t}V_{t+1}$ and $K(\beta_t)$, are unaffected by κ. Thus, a rise in κ only affects the arrival rate $\rho_t = \frac{1}{\kappa K(\beta_t)}$, slowing down the solution of the problem. Also, we see from the diagram that increasing the value of delay W_{t+1} increases β_t (slows the arrival rate ρ_t), raising the expected payoff $E^{\pi_t}V_{t+1}$ and the time devoted to the decision, $\kappa K(\beta_t)$. Similarly, a upward shift in the distribution of values $V_{i,t+1}$, in the sense of first-order stochastic dominance, will decrease β_t, since this raises the curve $g(\beta) \equiv \ln(E^\eta \exp(\beta)V_{t+1})$. Intuitively, a favorable shift in $\vec{V}_{t+1}$ makes the DM more eager to finish the problem, and less concerned about distinguishing correctly between the different options.

[11] In the continuous-time limit, the "immediate" solution occurs only when $E^\eta V_{t+\epsilon} \geq W_{t+\epsilon}$, for infinitesimal ϵ, and the solution to (19) reduces to $\beta_t = 0$, with instantaneous arrival.

2.2 Choosing a stopping time

The model in the previous section had three solution regimes (postponement, interior, or immediate solution), depending on how the value of delay, W_{t+1}, compared with the values $V_{i,t+1}$ available. While the model displayed error-prone decisions across the alternatives i, it could also be regarded as an *error-free* choice of the arrival rate ρ_t. That is, it could be regarded as a generalized optimal stopping problem, with the choice between stopping immediately ($\rho_t = 1$) and postponing ($\rho_t = 0$) supplemented by all intermediate probabilities $\rho_t \in (0, 1)$, in which the stopping hazard ρ_t was chosen without error. This observation raises the question of how we might allow for *errors in the stopping rate too*, for consistency with our treatment of the choice across the alternatives i.

Hence, we next study an error-prone stopping time problem, in which it is costly to choose the stopping rate ρ precisely.[12] Abstracting briefly from the choice across alternatives i, we assume the DM simply chooses between continuing a process, to obtain value W, or "stopping" that process to obtain value X. 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 $\bar{\rho}$. The constant $\bar{\rho}$ is a free parameter in our framework, relating to the underlying speed at which the DM is capable of making decisions. Cost increases with deviations from the actual hazard from the uniform hazard at rate $\bar{\rho}$, as measured by Kullback-Leibler divergence.[13] To make the problem meaningful, we also allow for an alternative use of time (besides decision-making), which we will call "work".

Let $h_t \in [0, 1]$ be time spent "working", and μ_t be time spent "monitoring" the situation, to decide whether or not to stop. Let the output from work be $f(h)$, where

$$f(h) \geq 0 \tag{20}$$

$$f'(h) > 0, \quad f''(h) < 0, \tag{21}$$

$$\lim_{h \to 0} f'(h) = \infty. \tag{22}$$

The stopping rate is governed by the following Bellman equation:

$$W_t = \max_{h,\mu,\rho} f(h) + (1-\delta)\left[\rho X_{t+1} + (1-\rho)W_{t+1}\right] \tag{23}$$

$$\text{subject to:} \quad \kappa \mathcal{D}\left(\begin{pmatrix} \rho \\ 1-\rho \end{pmatrix} \middle|\middle| \begin{pmatrix} \bar{\rho} \\ 1-\bar{\rho} \end{pmatrix}\right) \leq \mu, \tag{24}$$

$$\text{and} \quad h + \mu \leq 1. \tag{25}$$

The first constraint limits the precision of the choice of the stopping rate ρ. The relative entropy term in the control cost function is given by

$$\mathcal{D}\left(\begin{pmatrix} \rho \\ 1-\rho \end{pmatrix} \middle|\middle| \begin{pmatrix} \bar{\rho} \\ 1-\bar{\rho} \end{pmatrix}\right) = \rho \ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho)\ln\left(\frac{1-\rho}{1-\bar{\rho}}\right). \tag{26}$$

[12] The error-prone stopping decision in this subsection is analogous to the "Errors-in-timing" model considered in CN15.

[13] At first glance, one might wish to treat the choice to stop or continue in a single time step as a binary decision, applying the model from Sec. 2.1 with a uniform benchmark $\vec{u} \equiv \{0.5, 0.5\}$ on the two options. But such a model is not well behaved as we redefine time units; it would imply a 50% probability of adjustment in one period, conditional on indifference between stopping and continuation, *regardless of the length of the time period.* In other words, the continuous-time limit of such a model would imply perfectly rational timing, regardless of κ. The introduction of the benchmark stopping hazard $\bar{\rho}$ is what gives our model a well-behaved continuous-time limit in which timing errors are still possible.

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

$$\begin{aligned} f'(h) - \lambda_h &= 0 && (27)\\ \lambda_\mu - \lambda_h &= 0 && (28)\\ (1-\delta)(X_{t+1} - W_{t+1}) - \lambda_\mu \kappa \left(\ln\left(\frac{\rho}{\bar\rho}\right) + 1 - \ln\left(\frac{1-\rho}{1-\bar\rho}\right) - 1 \right) &= 0 && (29)\\ \mu - \kappa \mathcal{D} &= 0 && (30)\\ 1 - \mu - h &= 0. && (31) \end{aligned}$$

Eliminating the multipliers and rearranging, we obtain

$$\ln\left(\frac{\rho}{1-\rho}\right) = \ln\left(\frac{\bar\rho}{1-\bar\rho}\right) + \beta(X_{t+1} - W_{t+1}) \tag{32}$$

where the precision of the decision is indexed by β:

$$\beta = \frac{1-\delta}{\kappa f'(h)}. \tag{33}$$

Solving for the equilibrium stopping rate ρ_t, we obtain

$$\begin{aligned} \rho_t &= \frac{\bar\rho \exp(\beta_t X_{t+1})}{(1-\bar\rho)\exp(\beta_t W_{t+1}) + \bar\rho \exp(\beta_t X_{t+1})} && (34)\\ &= \frac{\bar\rho \exp(\beta_t D_{t+1})}{1 - \bar\rho + \bar\rho \exp(\beta_t D_{t+1})} && (35)\\ &= \frac{\bar\rho}{\bar\rho + (1-\bar\rho)\exp(-\beta_t D_{t+1})}, && (36) \end{aligned}$$

where $D_{t+1} = X_{t+1} - W_{t+1}$. This equation expresses the stopping probability ρ_t as a weighted binomial logit that places weight $\bar\rho$ on stopping, and weight $1-\bar\rho$ on continuation. Note that as the precision β_t approaches zero, the process stops with a constant hazard $\bar\rho$ regardless of the value of stopping, while as $\beta_t \to \infty$, the process stops if and only if $X_{t+1} > W_{t+1}$. The parameter $\bar\rho$ controls the *speed* of stopping; more exactly, it represents the stopping hazard when the DM is indifferent between stopping and continuing.Using this logit hazard formula, we can then write the decision entropy (26) as

$$\begin{aligned} \mathcal{D}\left(\begin{pmatrix} \rho_t \\ 1-\rho_t \end{pmatrix} \middle|\middle| \begin{pmatrix} \bar\rho \\ 1-\bar\rho \end{pmatrix} \right) &= \rho_t \beta_t D_{t+1} - \ln(1 - \bar\rho + \bar\rho \exp(\beta_t D_{t+1})) && (37)\\ &= \frac{\bar\rho \beta_t D_{t+1}}{\bar\rho + (1-\bar\rho)\exp(-\beta_t D_{t+1})} - \ln(1 - \bar\rho + \bar\rho \exp(\beta_t D_{t+1})) \equiv K_\rho(\beta_t). && (38) \end{aligned}$$

The function $K_\rho(\beta)$ represents a cost function for precision; differentiation shows (see proof of Prop. 2) that $K'_\rho(\beta) > 0$.

The model of Sec. 2.1 had two 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, which simplifies the analysis by ensuring an interior solution for β_t at any step of the backwards induction process defined by (23). To see this, write the time constraint as

$$1 - h = \mu = \kappa K_\rho(\beta), \tag{39}$$

where $K_\rho(\beta)$ is given by (38). Then we have:

Proposition 2 *Define $D_{t+1} \equiv X_{t+1} - W_{t+1}$, and suppose f satisfies (20)-(22). Then there is a unique pair (h_t, β_t), with $\beta_t > 0$ and $h_t \in (0, 1]$, that satisfy (33) and (39). The pair (h_t, β_t) solves the problem (23)-(25) conditional on X_{t+1} and W_{t+1}, and varies smoothly with D_{t+1}.*

Proof. Note that for $h \in [0, 1]$, the first-order condition (33) 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 $\beta_1^* \equiv \frac{1-\delta}{\kappa f'(1)}$ at $h = 1$. On the other hand, if we define $z \equiv \beta D_{t+1}$ and

$$\mathcal{D} \;\equiv\; \frac{\bar{\rho} z}{\bar{\rho} + (1 - \bar{\rho}) \exp(-z))} - \ln(1 - \bar{\rho} + \bar{\rho} \exp(z)) \;\equiv\; F_\rho(z).$$

then by differentiating and simplifying we obtain

$$F'_\rho(z) \;=\; \frac{z\bar{\rho} \exp(-z)}{[\bar{\rho} + (1 - \bar{\rho}) \exp(-z))]^2}.$$

Hence $F'_\rho(z) > (<) 0$ if and only if $z > (<) 0$, which requires $D_{t+1} > (<) 0$ (for $\beta > 0$). Therefore, writing $K'_\rho(\beta) = F'_\rho(z) \frac{\partial z}{\partial \beta}$, whenever $D_{t+1} > 0$, we conclude that $K'_\rho(\beta) > 0$ because $F'_\rho(z) > 0$ and $\frac{\partial z}{\partial \beta} = D_{t+1} > 0$. Likewise, when $D_{t+1} < 0$, we conclude that $K'_\rho(\beta) > 0$ because $F'_\rho(z) < 0$ and $\frac{\partial z}{\partial \beta} = D_{t+1} < 0$. If instead $D_{t+1} = 0$ exactly, then we can define $K_\rho(\beta) = 0$ for all $\beta < \infty$.

Therefore, if we graph (33) and (39) in (β, h) space, the curve (33) slopes upwards from the origin, and the curve (39) slopes downwards from $(\beta, h) = (0, 1)$, except in the limiting case $D_{t+1} = 0$, when (39) becomes the horizontal line $h = 1$. Hence there is a unique solution $\beta_t > 0$ and $h_t \in (0, 1]$, which is an interior solution unless $D_{t+1} = 0$ exactly. In the latter case, the solution reaches the corner $h_t = 1$, with $\rho_t = \bar{\rho}$, and we can define $\beta_t = \frac{1-\delta}{\kappa f'(1)}$, so that h_t, ρ_t, and β_t all vary continuously and smoothly with D_{t+1}. **QED.**

Like Prop. 1, Prop. 2 shows that there is a unique solution to a single iterative step of the Bellman equation, given the forward values X_{t+1} and W_{t+1}. Graphed with β on the horizontal axis and h on the vertical axis, (33) slopes up, and (39) slopes down, with exactly one interior crossing under the conditions stated in the proposition. A rise in κ shifts both curves left, decreasing the optimal β_t and thus bringing ρ_t closer to $\bar{\rho}$. When $D_{t+1} > 0$, an increase in D_{t+1} shifts (39) left, decreasing β_t while ρ_t rises further above $\bar{\rho}$. If instead $D_{t+1} < 0$, an increase in D_{t+1} shifts (39) right, increasing β_t while raising ρ_t towards $\bar{\rho}$.

2.3 Choosing when to solve a problem

A situation in which a DM can choose between several options, but may also simply postpone that choice, effectively combines a static decision problem with an optimal stopping problem. A game, such as the bargaining game we consider later, may include both of these elements at various stages.

Therefore, we now study a situation that combines the two decisions analyzed thus far: a DM, subject to control costs, chooses a stopping time that determines when he will draw a conclusion between several options. As in Section 2.1, the DM chooses a solution rate ρ for the choice problem, and a quicker solution implies a less precise choice across the options considered. But instead of assuming that ρ is chosen optimally, here we also constrain the precision of the choice of the solution rate ρ, as in Sec. 2.2. Errors in the choice of ρ amount to errors in the timing of the choice, as distinct from an error in which option is chosen.[14]

[14] The model developed in this subsection is analogous to the benchmark specification in CN15.

Figure 3: Choosing when to solve a problem.

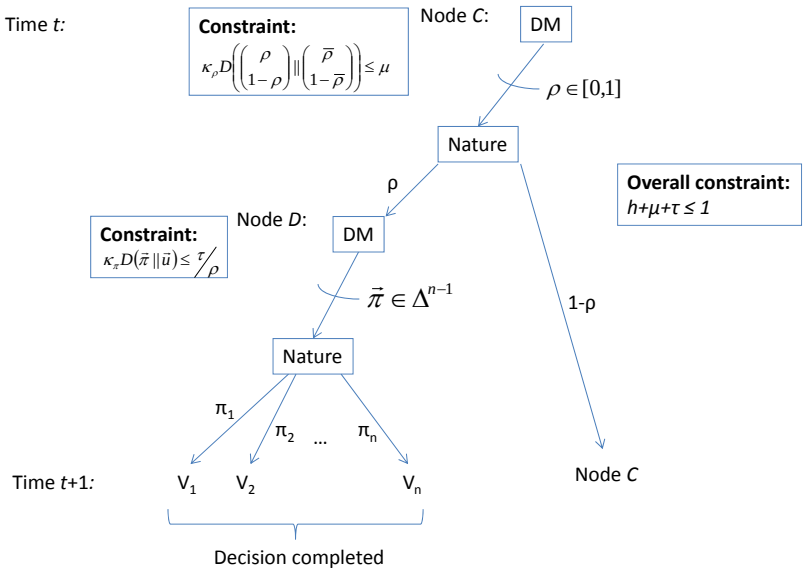

Note: Decision maker (DM) chooses solution rate ρ_t of decision; Nature lets conclusion arrive with probability ρ_t. A slower decision allows more precise allocation of probabilities $\pi_{i,t}$ across alternatives i with values $V_{i,t+1}$, satisfying the constraint $\kappa_\pi \mathcal{D}(\vec{\pi}_t||\vec{\eta}) \leq \tau_t/\rho_t$, where τ_t is the fraction of time spent comparing alternatives i. Choosing the solution rate ρ_t is also costly, and must satisfy the constraint (42), where μ_t is the fraction of time devoted to choosing ρ_t. Time subscripts suppressed in diagram.

This now leads us to consider three possible uses of time. First, let $h \in [0,1]$ be the time dedicated to non-decision activities, such as work or leisure. As before, we assume that non-decision activities yield payoffs $f(h)$, satisfying (20)-(22). Second, let $\mu \in [0,1]$ be the fraction of time spent "monitoring" whether the current moment is a good time to solve the decision problem. The precision of the solution rate ρ will be limited by monitoring time, under a constraint identical to (24). Finally, let $\tau \in [0,1]$ be the fraction of time dedicated to actually choosing between options i. Since only a fraction of the DM's time is dedicated this choice, the precision of this choice problem is no longer constrained by $1/\rho$ (the expected duration of the decision), as in (1). Instead, precision will be constrained by τ/ρ (expected time devoted to the decision problem, up until it is solved). Therefore, we now describe the decision process by the

following Bellman equation:

$$W_t = \max_{\rho, h, \tau, \mu, \{\pi_i\}_{i=1}^n} f(h) + (1-\delta)\left[\rho \sum_{i=1}^n \pi_i V_{i,t+1} + (1-\rho) W_{t+1}\right] \tag{40}$$

$$\text{s.t.:} \quad \rho\kappa_\pi \sum_{i=1}^n \pi_i \ln\left(\frac{\pi_i}{\eta_i}\right) \leq \tau, \tag{41}$$

$$\text{and} \quad \kappa_\rho\left(\rho \ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho)\ln\left(\frac{1-\rho}{1-\bar{\rho}}\right)\right) \leq \mu, \tag{42}$$

$$\text{and} \quad \sum_{i=1}^n \pi_i = 1, \quad \text{and} \quad h + \mu + \tau \leq 1. \tag{43}$$

Here, we have allowed for the possibility that precision in choice and precision in timing may have different time costs, paramerized by κ_π and κ_ρ, respectively.

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:

$$\lambda_h = f'(h) = \lambda_\tau = \lambda_\mu. \tag{44}$$

The remaining first-order conditions can be written as

$$(1-\delta)\rho V_{i,t+1} - \lambda_\tau \rho \kappa_\pi \left(1 + \ln\left(\frac{\pi_i}{\eta_i}\right)\right) - \lambda_\pi = 0 \tag{45}$$

$$(1-\delta)\left(\sum_{i=1}^n \pi_i V_{i,t+1} - W_{t+1}\right) - \lambda_\tau \kappa_\pi \mathcal{D}_\pi - \lambda_\mu \kappa_\rho \left[\ln\left(\frac{\rho}{\bar{\rho}}\right) + 1 - \ln\left(\frac{1-\rho}{1-\bar{\rho}}\right) - 1\right] = 0 \tag{46}$$

$$\tau - \rho\kappa_\pi \mathcal{D}_\pi = 0 \tag{47}$$

$$\mu - \kappa_\rho \mathcal{D}_\rho = 0 \tag{48}$$

$$1 - h - \mu - \tau = 0 \tag{49}$$

$$1 - \sum_{i=1}^n \pi_i = 0 \tag{50}$$

where

$$\mathcal{D}_\pi \equiv \sum_{i=1}^n \pi_i \ln\left(\frac{\pi_i}{\eta_i}\right) \tag{51}$$

$$\mathcal{D}_\rho \equiv \rho \ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho)\ln\left(\frac{1-\rho}{1-\bar{\rho}}\right). \tag{52}$$

As before, (45) and (50) imply that the optimal probabilities $\pi_{i,t}$ take the form of a logit:

$$\pi_{i,t} = \frac{\eta_i \exp(\beta_{\pi,t} V_{i,t+1})}{\sum_j \eta_j \exp(\beta_{\pi,t} V_{j,t+1})}, \tag{53}$$

where

$$\beta_{\pi,t} = \frac{1-\delta}{\kappa_\pi f'(h_t)}. \tag{54}$$

Following equations (12)-(13), we can calculate the entropy cost $\mathcal{D}_{\pi,t}$ paid to choose $\pi_{i,t}$:

$$\mathcal{D}_{\pi,t} \;=\; \beta_{\pi;t}E^{\pi_t}V_{t+1} - \ln(E^{\eta}\exp(\beta_{\pi,t}V_{t+1})) \;=\; \beta_{\pi,t}\frac{E^{\eta}V_{t+1}\exp(\beta_{\pi,t}V_{t+1})}{E^{\eta}\exp(\beta_{\pi,t}V_{t+1})} - \ln(E^{\eta}\exp(\beta_{\pi,t}V_{t+1})). \tag{55}$$

Following (32)-(36), we can again show that the solution rate ρ_t takes logit form:

$$\rho_t \;=\; \frac{\bar{\rho}}{\bar{\rho} + (1-\bar{\rho})\exp\left(-\beta_{\rho,t}D_{t+1}\right)} \tag{56}$$

$$\;=\; \frac{\bar{\rho}\exp\left(\beta_{\rho,t}\tilde{V}_{t+1}\right)}{\bar{\rho}\exp\left(\beta_{\rho,t}\tilde{V}_{t+1}\right) + (1-\bar{\rho})\exp\left(\beta_{\rho}W_{t+1}\right)} \tag{57}$$

$$\;=\; \frac{\bar{\rho}\left(E^{\eta}\exp(\beta_{\pi,t}V_{t+1})\right)^{\frac{\kappa_\pi}{\kappa_\rho}}}{\bar{\rho}\left(E^{\eta}\exp(\beta_{\pi,t}V_{t+1})\right)^{\frac{\kappa_\pi}{\kappa_\rho}} + (1-\bar{\rho})\left(\exp\left(\beta_{\pi,t}W_{t+1}\right)\right)^{\frac{\kappa_\pi}{\kappa_\rho}}}, \tag{58}$$

where

$$\beta_{\rho,t} \;\equiv\; \frac{1-\delta}{\kappa_\rho f'(h_t)} \;=\; \frac{\kappa_\pi}{\kappa_\rho}\beta_{\pi,t}, \tag{59}$$

$$D_{t+1} \;\equiv\; \tilde{V}_{t+1} - W_{t+1}, \tag{60}$$

$$\tilde{V}_{t+1} \;\equiv\; E^{\pi_t}V_{t+1} - \frac{\mathcal{D}_{\pi,t}}{\beta_{\pi,t}}. \tag{61}$$

Equation (57) shows that the solution rate is given by a weighted binary logit comparing the values of decision and postponement; it corresponds to equation (20) of CN15.

Equation (59) 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 $\tilde{V}_{t+1}$ represents the expected gains that accrue upon adjustment, net of adjustment costs; the factor $\beta_{\pi,t}$ converts precision into time at rate κ_π, time into utility at rate $f'(h_t)$, and next period's utility into utility now at rate $1-\delta$, so that the payoff $E^{\pi_t}V_{t+1}$ is commensurate with the precision measure $\mathcal{D}_{\pi,t}$. Using (55), we can calculate

$$\tilde{V}_{t+1} \;=\; \beta_{\pi,t}^{-1}\ln\left(E^{\eta}\exp\left(\beta_{\pi,t}V_{t+1}\right)\right). \tag{62}$$

This corresponds to the explicit formula (15) for the value function that is stated in CN15. Finally, following (37)-(38), the "monitoring" time devoted to choosing the solution rate is

$$\mathcal{D}_{\rho,t} \;=\; \frac{\bar{\rho}\beta_{\rho,t}D_{t+1}}{\bar{\rho} + (1-\bar{\rho})\exp\left(-\beta_{\rho,t}D_{t+1}\right)} - \ln(1-\bar{\rho}+\bar{\rho}\exp(\beta_{\rho,t}D_{t+1})). \tag{63}$$

Do the first-order conditions (53)-(61) define a unique solution for problem (40)-(43)? To address this question, note that conditional on the values $\vec{V}_{t+1}$ and W_{t+1} associated with the time $t+1$ 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:

$$1-h \;=\; \rho\kappa_\pi\mathcal{D}_\pi + \kappa_\rho\mathcal{D}_\rho, \tag{64}$$

$$\kappa_\pi\beta_\pi f'(h) \;=\; 1-\delta. \tag{65}$$

These two equations (illustrated by Fig. 4) effectively depend on only two variables, h and β_π, because all other time-t endogenous variables are easily substituted out. $\mathcal{D}_\pi$ can be eliminated using equation (55), and ρ can be eliminated using (56) or (58). If we also calculate β_ρ and $\tilde{V}_{t+1}$ using (59) and (62), we can then eliminate $\mathcal{D}_\rho$ using (63).

We can now show that the two equations (64)-(65) have a unique solution. To do so, it helps to distinguish the special case $\kappa_\pi = \kappa_\rho$ from the general case in which these parameters differ.

2.3.1 Special case: Unifying timing and choice across alternatives when $\kappa_\pi = \kappa_\rho$

Consider the special case $\kappa_\pi = \kappa_\lambda \equiv \kappa$. In this case, it is helpful to rewrite the decision problem in terms of the following set of alternatives:

$$\Gamma^\dagger \;\equiv\; \{\text{Postponement}\} \cup \Gamma.$$

This choice situation is equivalent to that modeled in problem (40), but simply treats postponement as one more alternative to be considered, rather than classifying it separately from the other options. Without loss of generality, suppose the firm chooses postponement with probability $1-\rho$, and draws from the remaining alternatives with probabilities $\pi_i^\dagger = \rho\pi_i$, where $1 - \rho + \sum_{i=1}^n \pi_i^\dagger = 1$.

Now, using the parameter restriction $\kappa_\pi = \kappa_\lambda$, the two cost functions (41)-(42) can be summed as follows:

$$\begin{aligned}
\rho\kappa_\pi\mathcal{D}_\pi + \kappa_\rho\mathcal{D}_\rho \;&= \rho\kappa_\pi \sum_i \pi_i \ln\left(\frac{\pi_i}{\eta_i}\right) \;+\; \kappa_\rho\left[\rho\ln\left(\frac{\rho}{\bar\rho}\right) + (1-\rho)\ln\left(\frac{1-\rho}{1-\bar\rho}\right)\right] \\
&= \kappa\left\{(1-\rho)\ln\left(\frac{1-\rho}{1-\bar\rho}\right) \;+\; \rho\left[\sum_i \pi_i \ln\left(\frac{\pi_i}{\eta_i}\right) \;+\; \ln\left(\frac{\rho}{\bar\rho}\right)\right]\right\} \qquad (66) \\
&= \kappa\left\{(1-\rho)\ln\left(\frac{1-\rho}{1-\bar\rho}\right) \;+\; \rho\sum_i \pi_i \ln\left(\frac{\rho\pi_i}{\bar\rho\eta_i}\right)\right\} \\
&= \kappa\left\{(1-\rho)\ln\left(\frac{1-\rho}{1-\bar\rho}\right) \;+\; \sum_i \pi_i^\dagger \ln\left(\frac{\pi_i^\dagger}{\bar\rho\eta_i}\right)\right\}. \qquad (67)
\end{aligned}$$

The expression in (67) is a relative entropy measure over the set of alternatives $\Gamma_t^\dagger$: it is the Kullback-Leibler divergence of the probabilities $(1-\rho, \vec{\pi}^\dagger) \equiv (1-\rho, \pi_1^\dagger, \ldots \pi_n^\dagger)$ from the default distribution $(1-\bar\rho, \bar\rho\eta_1, \ldots, \bar\rho\eta_n)$. It is a convex function of the probabilities $(1-\rho, \vec{\pi}^\dagger)$.

Therefore, when $\kappa_\rho = \kappa_\pi \equiv \kappa$, problem (40)-(43) can be rewritten equivalently as follows:

$$W_t \;=\; \max_{\rho, h, \tau^\dagger, \{\pi_i^\dagger\}_{i=1}^n} \; f(h) + (1-\delta)\left[\rho\sum_{i=1}^n \pi_i^\dagger V_{i,t+1} + (1-\rho)W_{t+1}\right] \qquad (68)$$

$$\text{s.t.:} \qquad \kappa\left((1-\rho)\ln\left(\frac{1-\rho}{1-\bar\rho}\right) + \sum_{i=1}^n \pi_i^\dagger \ln\left(\frac{\pi_i^\dagger}{\bar\rho\eta_i}\right)\right) \;\le\; \tau^\dagger, \qquad (69)$$

$$\text{and} \qquad \sum_{i=1}^n \pi_i^\dagger = \rho, \quad \text{and} \quad h + \tau^\dagger \le 1, \qquad (70)$$

where $\tau^\dagger \equiv \tau + \mu$ is total time used for decision-making (including "monitoring"). Note that the objective function in (68) is equivalent to that in (40); and if $\kappa_\rho = \kappa_\pi \equiv \kappa$, then the constraints (69)-(70) are satisfied if and only (41)-(43) hold.

Thus, the firm's choice can be regarded as a single decision across the set of alternatives $\Gamma^\dagger$, subject to a relative entropy constraint. An advantage of writing the problem this way is that it maximizes a concave function over a convex set. Hence we can show that each backwards induction step has a unique solution:

Proposition 3 *Suppose f satisfies (20)-(22) for $h \in [0,1]$, and $\min_i V_{i,t+1} < \max_i V_{i,t+1}$. Then problem (68)-(70) has a unique solution $h_t \in (0,1)$, $\beta_t > 0$, which is characterized by the first-order conditions (64)-(65).*

If $\kappa_\pi = \kappa_\rho = \kappa$, this also represents a solution of the original problem (40)-(43).

Proof. Since the payoff function is linear in the probabilities $(1-\rho, \vec{\pi}^\dagger)$, and concave in h, it is concave overall. Since relative entropy is a strictly convex function, (69) defines a convex set, as do the simplex constraints in (70). Thus, the Bellman equation maximizes a stictly concave function over a convex set (the intersection of the constraint sets), so there exists a unique solution to the problem.

Now, if we graph (65) in (β, h) space, assumptions (20)-(22) imply that it slopes upwards from the origin. Plugging (55), (58), and (63) into (67), we see that the total decision cost is zero when $\beta = 0$, and is strictly positive for all $\beta > 0$ unless $\min_i V_{i,t+1} = \max_i V_{i,t+1} = W_{t+1}$. Therefore curve (64) passes through the point $(\beta, h) = (0,1)$ and lies strictly below $h = 1$ for all $\beta > 0$ if $\min_i V_{i,t+1} < \max_i V_{i,t+1}$. Therefore (64) and (65) have a solution $h \in (0,1)$, $\beta \in (0, \frac{1-\delta}{\kappa f'(1)})$ which represents an interior solution for the problem (68)-(70). **QED.**

2.3.2 General case

For the general case where $\kappa_\pi \neq \kappa_\rho$, we can show the existence of a unique solution to any backwards induction step by showing that the slopes of (64) and (65) differ in sign. As Fig. 4 shows, first-order condition (65) gives h as an increasing function of β_π. It is therefore helpful to analyze the slope of the time constraint (64), viewed as a relation between β_π and h. It will be helpful to use the identities $\beta_\pi D_{t+1} = \ln(E^\eta \exp(\beta_\pi(V - W_{t+1})))$ and $\exp\left(\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+1}\right) = (E^\eta \exp(\beta_\pi(V - W_{t+1})))^{\kappa_\pi/\kappa_\rho}$. Substituting, the time constraint can be written as follows:

$$\begin{aligned}
1 - h &= \rho(\beta_\pi)\kappa_\pi \left(\beta_\pi \frac{E^\eta V e^{\beta_\pi V}}{E^\eta e^{\beta_\pi V}} - \ln(E^\eta e^{\beta_\pi V})\right) + \kappa_\rho \left[\rho(\beta_\pi)\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+1} - \ln\left(1 - \bar{\rho} + \bar{\rho} e^{\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+1}}\right)\right] \\
&= \kappa_\pi \beta_\pi \rho(\beta_\pi)\left(\frac{E^\eta V e^{\beta_\pi V}}{E^\eta e^{\beta_\pi V}}\right) - \kappa_\pi \beta_\pi \rho(\beta_\pi) W_{t+1} - \kappa_\rho \ln\left(1 - \bar{\rho} + \bar{\rho} e^{\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+1}}\right) \\
&= \kappa_\pi \beta_\pi \rho(\beta_\pi)\left[\frac{E^\eta V e^{\beta_\pi V}}{E^\eta e^{\beta_\pi V}} - W_{t+1}\right] - \kappa_\rho \ln\left(1 - \bar{\rho} + \bar{\rho} e^{\frac{\kappa_\pi}{\kappa_\rho}\beta_\pi D_{t+1}}\right) \;\equiv\; t(\beta_\pi), \qquad (71)
\end{aligned}$$

where $\rho(\beta_\pi)$ is the function defined by (58). Note that, conditional on $\vec{V}_{t+1}$ and W_{t+1}, the right-hand side of (71) 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 (64) slopes downward, and hence the model has a unique interior optimum.

Figure 4: Time use in equilibrium.

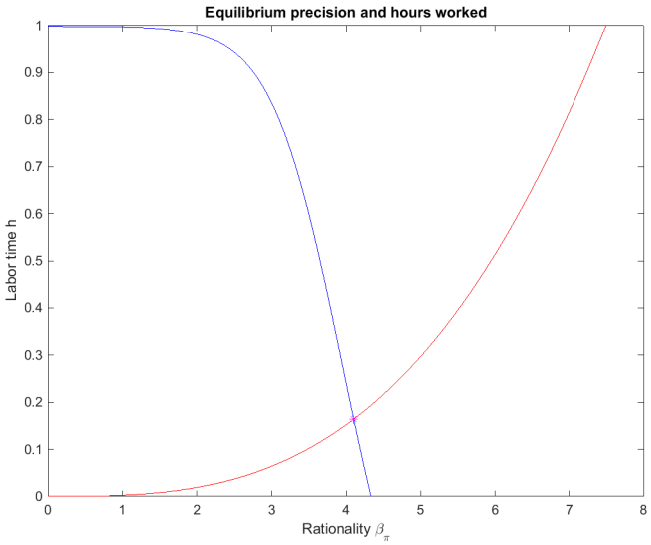

Note: Choice of precision β_π and labor time h.
First-order condition (65) in red, and time constraint (64) in blue.
Parameters: $\delta = 0.001$, $\bar{\rho} = 0.01$, $\kappa_\pi = \kappa_\rho = 0.2$, $f(h) = Zh^\zeta$ with $Z = 1$, $\zeta = 2/3$. Also $W_{t+1} = 9$, and $\vec{V}_{t+1} = 11 - (5 - \vec{x})^2$, where $\vec{x} = [-10, -9, \ldots, 9, 10]$.

Proposition 4 *Suppose f satisfies (20)-(21) for $h \in [0,1]$. Then problem (40)-(43) is solved by the pair $h_t \in (0,1]$, $\beta_{\pi,t} > 0$ that correspond to the unique crossing of curves (64)-(65). The solution $(h_t, \beta_{\pi,t})$ varies smoothly with changes in $\vec{V}_{t+1}$ and W_{t+1}. The corner solution $h_t = 1$ arises only when $\min_i V_{i,t+1} = \max_i V_{i,t+1} = W_{t+1}$; otherwise $h_t < 1$ strictly.*[15]

Proof. The main point to prove is that the right-hand side of (71) is nondecreasing in β. Let us write (71) as $h = 1 - t(\beta)$, where

$$t(\beta) \;=\; \kappa_\pi \beta \rho(\beta) \left(\frac{m'(\beta)}{m(\beta)} - W \right) - \kappa_\rho \ln \left(1 - \bar{\rho} + \bar{\rho} \left(\frac{m(\beta)}{e^{\beta W}} \right)^{\frac{\kappa_\pi}{\kappa_\rho}} \right), \tag{72}$$

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

[15] A MATLAB program, `twomargins_eta.m`, is provided to solve the problem analyzed in Proposition 4, taking as given the values W_{t+1} and $\vec{V}_{t+1}$, the default probabilities $\vec{\eta}$, and other parameters. The crossing of curves (64)-(65) is calculated by bisection.

Now by differentiating and simplifying, we can show that

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

using (58). Notice that the first and last terms cancel; the remaining terms are

$$t'(\beta) = \frac{\kappa_\pi^2}{\kappa_\rho} \rho(\beta) \left(\frac{m'}{m} - W\right)^2 + \kappa_\pi \beta \rho \left(\frac{m''}{m} - \left(\frac{m'}{m}\right)^2\right). \tag{73}$$

The first term on the right-hand side of (73) is obviously nonnegative. In the second term on the right-hand side of (73), we have

$$\frac{m''}{m} - \left(\frac{m'}{m}\right)^2 = \frac{E^\eta V^2 \exp(\beta V)}{E^\eta \exp(\beta V)} - \left(\frac{E^\eta V \exp(\beta V)}{E^\eta \exp(\beta V)}\right)^2 = E^\pi V^2 - (E^\pi V)^2 \geq 0. \tag{74}$$

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 (73) 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 $h = 1 - t(\beta_\pi)$ is weakly downward-sloping from $h = 1$ at $\beta_\pi = 0$. It is strictly downward-sloping except in the limiting case $\min_i V_{t+1} = \max_i V_{t+1} = W_{t+1}$ This curve shifts smoothly with any change in $\vec{V}_{t+1}$ or W_{t+1}. The curve $f'(h) = \frac{1-\delta}{\kappa_\pi \beta_\pi}$ slopes upward from $\beta_0 \equiv \frac{1-\delta}{\kappa_\pi f'(0)} = 0$; it exceeds $h = 1$ for $\beta_\pi > \frac{1-\delta}{\kappa_\pi f'(1)}$.[16] Hence the two curves have a interior unique crossing which varies smoothly with changes in $\vec{V}_{t+1}$ and W_{t+1}. **QED.**

2.3.3 Backwards induction

Propositions 3-4 established general conditions under which a backwards induction step of the form (40)-(43) has a unique, well-behaved optimum when future values $\vec{V}_{t+1}$ and W_{t+1} are taken as given. Now, we show that each backwards induction step acts as a contraction, implying that (40)-(47) also has a unique, well-behaved optimum when interpreted as an infinite-horizon dynamic programming problem.

Proposition 5 *Suppose f satisfies (20)-(21). Then:*
(a.) the operator defined by problem (40)-(43) is a contraction, and
(b.) if the problem is defined over an infinite horizon and the options $i = 1..n$ have time-independent values $\vec{V}_{ss}$ then (40)-(47) has a unique time-independent value $W_t = W_{t+1} \equiv W_{ss}$.

[16] The Inada condition (22) is not necessary to prove Prop. 4. It suffices to assume that $\lim_{h \to 0} f'(h) > \frac{\kappa_\pi}{1-\delta}$.

Proof. Consider two possible continuation values W^a_{t+1} and W^b_{t+1}, where $W^a_{t+1} = W^b_{t+1} - 1$. Let $(h^a, \rho^a, \vec{\pi}^a)$ be the policy that solves (40) when $W_{t+1} = W^a_{t+1}$, with time t value W^a_t; and let $(h^b, \rho^b, \vec{\pi}^b)$ be the solution when $W_{t+1} = W^b_{t+1} \equiv W^a_{t+1} + 1$, implying time t value W^b_t. Then we have

$$
\begin{aligned}
W^+_t &\equiv f(h^b) + (1-\delta)W^a_{t+1} + \rho^b(1-\delta)(E^{\pi^b}V - W^a_{t+1}) + (1-\delta) && (75)\\
&\geq f(h^b) + (1-\delta)W^a_{t+1} + \rho^b(1-\delta)(E^{\pi^b}V - W^a_{t+1}) + (1-\rho^b)(1-\delta) && (76)\\
= W^b_t &= f(h^b) + (1-\delta)W^b_{t+1} + \rho^b(1-\delta)(E^{\pi^b}V - W^b_{t+1}) && (77)\\
&\geq f(h^a) + (1-\delta)W^b_{t+1} + \rho^a(1-\delta)(E^{\pi^a}V - W^b_{t+1}) && (78)\\
&= f(h^a) + (1-\delta)W^a_{t+1} + \rho^a(1-\delta)(E^{\pi^a}V - W^a_{t+1}) + (1-\rho^a)(1-\delta) && (79)\\
\geq W^a_t &= f(h^a) + (1-\delta)W^a_{t+1} + \rho^a(1-\delta)(E^{\pi^a}V - W^a_{t+1}) && (80)\\
\geq W^-_t &= f(h^b) + (1-\delta)W^a_{t+1} + \rho^b(1-\delta)(E^{\pi^b}V - W^a_{t+1}). && (81)
\end{aligned}
$$

Note that

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

whereas

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

Thus (40) defines a contraction of modulus $(1-\delta)$. Part (b) is an immediate corollary. **QED**.

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 of size K. Each may propose a share between 0 and 1 to the other player, and may accept or reject an offer received from the other player. Accepting the other's offer 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. 4.

We compare four possible bargaining protocols, representing different assumptions about the set of signals that the players may exchange over the course of play.

- One particularly simple protocol is the **rejectable offers protocol** where the possible responses to any offer are signals meaning "yes" or "no". Accepting the offer by saying "yes" ends the game, while rejection returns the game to a situation with no offer outstanding. The player who made a given offer cannot revise or withdraw that offer.

- A simple alternative is the **alternating offers protocol** where acceptance ("yes") ends the game, but rejection is equated with making a counteroffer. Therefore the responding player cannot "just say no", but must instead propose an alternative. Again, the proposer of a given offer cannot alter it.

- The **withdrawable offers protocol** generalizes the rejectable offers protocol by allowing *either* player to say "no" to the current offer, returning the game to a situation with no offer outstanding. That is, the recipient of an offer can accept it or reject it without offering an alternative, and the player that made the offer can likewise withdraw it without offering an alternative.

- Finally, the **updateable offers protocol** generalizes the alternating offers protocol by allowing *either* players to propose an alternative to the current offer. Thus, the responding player may accept or make a counterproposal, while the player that made a given proposal may instead substitute it by an updated proposal.

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 might be expected to imply substantial delay, compared with the other specifications. The updateable and withdrawable protocols are of particular interest because they are especially robust versions of the game, since they never oblige players to remain in an undesirable state.

We will solve the games by backwards induction from a finite ending time. Time steps are assumed discrete but very short, so we ignore the possibility that offers arrive simultaneously. Therefore in any of the four protocols considered, each player $i \in \{A, B\}$ may be in one of three types of states at any time during play:

- 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, player A's state is N_t^A if and only if B's state is N_t^B. The state of player i is $M_t^i(s)$ if and only if the other player, called $-i$, is in state $R_t^{-i}(s)$. Without loss of generality, we define offers from the point of view of agent A. Thus $s \in \Gamma^s \subset [0, 1]$ represents A's share of the cake, where Γ^s is a finite grid of shares between zero and one. Terminal payoffs are $u^A(s) \equiv sK$ and $u^B(s) \equiv (1-s)K$ for A and B, respectively. To avoid unnecessary notation, we will identify the names of the states with the names of the value functions associated with those states.

3.1 Solving the rejectable offers game

First, we analyze a bargaining game under the rejectable offers protocol. Figures 5-7 illustrate the nodes N_t^i, $R_t^i(s)$, and $M_t^i(s)$ under this protocol. We have used the notation ρ_t^i for player i's decision rate at node N_t^i, and $\pi_t^i(s)$ indicates the probabilities of the possible offers $s \in \Gamma^s$ which could be made at that node. Fig. 5 shows that players i and $-i$ choose their decision rates ρ_t^i and ρ_t^{-i} simultaneously (the figure is drawn from i's point of view; it could be redrawn from the perspective of $-i$ by relabelling nodes appropriately).[17] If player i completes a decision, the game will be distributed across nodes $M_{t+1}^i(s)$ at time $t+1$: 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 $R_{t+1}^i(s)$, indicating that an offer s has been made by $-i$.[18]

[17] The simultaneous decision is indicated by the dashed oval surrounding the subnodes DM_t^{-i}, which represents an information set. This indicates that all instances of the subnode DM_t^{-i} lie in the same information set, meaning that player $-i$ does not observe the value of player i's choice ρ_t^i. Since the figure ignores simultaneous arrivals, the time step should be chosen short enough so that the resulting inaccuracy is small.

[18] To conserve space, Fig. 5 suppresses the intermediate subnodes (shown in Fig. 3) at which Nature selects which of the decision outcomes will actually occur, after players i and $-i$ choose the probability allocations $\pi_t^i(s)$ and $\pi_t^{-i}(s)$.

Figure 5: Node N_t^i: no offer outstanding.

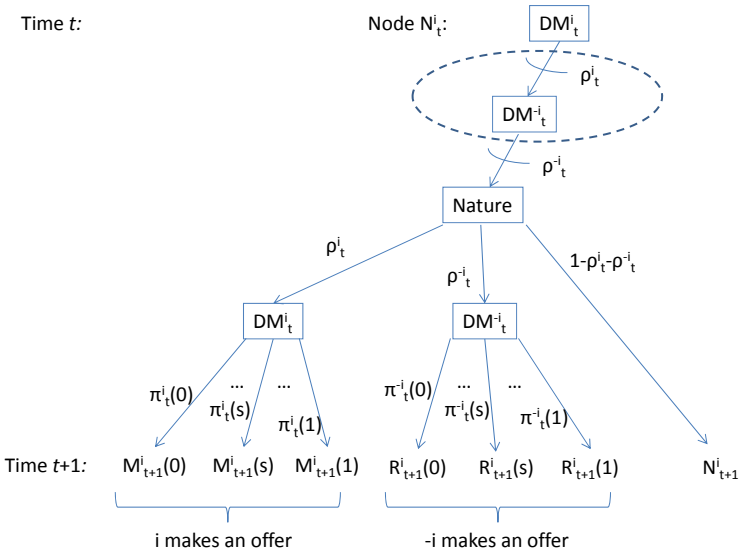

Note: Decision makers A and B simultaneously choose arrival rates ρ_t^A and ρ_t^B of their offers. A slower decision implies higher precision in the allocation of probabilities $\pi_t^i(s)$ across possible offers $s \in [0,1]$.

To describe the decision problems in the game, it is easiest to work backwards from node $R_t^i(s)$, at which player i is considering offer s received from $-i$, as illustrated in Figure 6. In this figure, $\lambda_t^i(s)$ represents the arrival rate of player i's decision, to accept with probability $\alpha_t^i(s)$ or reject with probability $1-\alpha_t^i(s)$. Policies $\lambda_t^i(s)$ and $\alpha_t^i(s)$ solve the following problem:

$$R_t^i(s) = \max_{\lambda,h,\tau,\mu,\alpha} f(h) + (1-\delta)\left\{\lambda\left[(1-\alpha)N_{t+1}^i + \alpha u^i(s)\right] + (1-\lambda)R_{t+1}^i\right\} \quad (82)$$

$$\text{s.t.:} \quad \lambda\kappa_\alpha\left(\alpha\ln\left(\frac{\alpha}{\bar{\alpha}}\right) + (1-\alpha)\ln\left(\frac{1-\alpha}{1-\bar{\alpha}}\right)\right) \leq \tau, \quad (83)$$

$$\text{and} \quad \kappa_\lambda\left(\lambda\ln\left(\frac{\lambda}{\bar{\lambda}}\right) + (1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right)\right) \leq \mu, \quad \text{and} \quad h+\mu+\tau \leq 1. \quad (84)$$

The Bellman equation (82)-(84) is formally identical to (40)-(43), so Prop. 4 applies, allowing us to calculate the policies and the value $R_t^i(s)$ in terms of the future values $R_{t+1}^i(s)$ and N_{t+1}^i.

Knowing the solution at node $R_t^{-i}(s)$, we can also calculate the value at node $M_t^i(s)$, where player i is waiting for a response from $-i$, as seen in Figure 7. Here there is no decision to be made; player i simply waits for a response, setting $h=1$. Therefore $M_t^i(s)$ satisfies:

$$M_t^i(s) = f(1) + (1-\delta)\left\{\lambda_t^{-i}(s)\left[(1-\alpha_t^{-i}(s))N_{t+1}^i + \alpha_t^{-i}(s)u^i(s)\right] + (1-\lambda_t^{-i}(s))M_{t+1}^i(s)\right\}. \quad (85)$$

The subgame at a node where no offer has yet been made is slightly more complex, because, even taking as given the time $t+1$ equilibrium, A's choice at time t interacts with B's choice at t. Since we have assumed shares are chosen from the grid Γ^s, the values that i chooses from

Figure 6: Rejectable offers protocol, node $R_t^i(s)$: offer of player $-i$ outstanding.

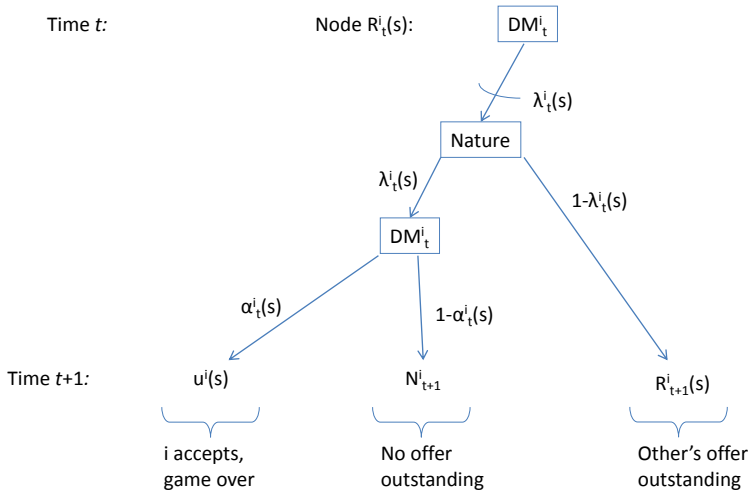

Note: Player i chooses response arrival rate λ_t^i. A slower response implies higher precision in the decision to accept or reject the offer s.

Figure 7: Rejectable offers protocol, node $M_t^i(s)$: offer of player i outstanding.

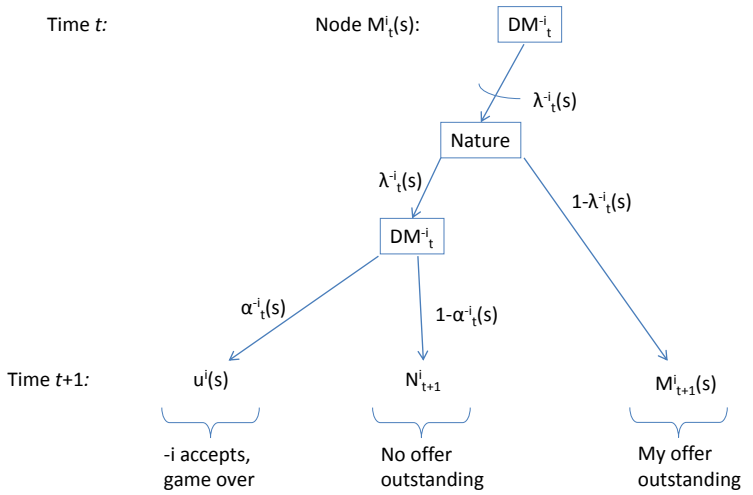

Note: Player $-i$ chooses response arrival rate λ_t^{-i}. A slower response implies higher precision in the decision to accept or reject the offer s.

at node N_t^i are $\vec{V}_{t+1} \equiv \{M_{t+1}^i(s) : s \in \Gamma^s\}$. If instead player i fails to make a decision at node N_t^i, he will face an offer from $-i$ with probability ρ_t^{-i}, or will return to a node with no offer outstanding with probability $1 - \rho_t^{-i}$. Therefore the continuation value for player i's problem is

$$W_{t+1}^i \;=\; (1 - \rho_t^{-i}) N_{t+1}^i + \rho_t^{-i} \sum_{s \in \Gamma^s} \pi_t^{-i}(s) R_{t+1}^i(s). \tag{86}$$

But the continuation value W_{t+1}^i is unknown unless we know the decision rate ρ_t^{-i} and probability allocation $\pi_t^{-i}(s)$ of player $-i$; hence we must solve the problems of the two players simultaneously. We can find a Nash equilibrium iteratively by beginning with the initial guess $W_{t+1}^i \approx N_{t+1}^i$, and then solving the following problem, using Prop. 4.

$$N_t^i \;=\; \max_{\rho, h, \tau, \mu, \{\pi(s)\}_{s \in \Gamma^s}} f(h) + (1 - \delta) \left[\rho \sum_{s \in \Gamma^s} \pi^i(s) M_{t+1}^i(s) + (1 - \rho) W_{t+1} \right] \tag{87}$$

$$\text{s.t.:} \qquad \rho \kappa_\pi \sum_{s \in \Gamma^s} \pi^i(s) \ln \left(\frac{\pi^i(s)}{\eta(s)} \right) \;\leq\; \tau, \tag{88}$$

$$\text{and} \qquad \kappa_\rho \left(\rho \ln \left(\frac{\rho}{\bar{\rho}} \right) + (1 - \rho) \ln \left(\frac{1 - \rho}{1 - \bar{\rho}} \right) \right) \;\leq\; \mu, \tag{89}$$

$$\text{and} \qquad \sum_{s \in \Gamma^s} \pi^i(s) = 1, \quad \text{and} \quad h + \mu + \tau \leq 1. \tag{90}$$

Solving this problem for both i and $-i$, we can update our guesses for W_{t+1}^i and W_{t+1}^{-i} using (86), and then solve (87)-(90) again for both players. We continue iterating in this way until we find fixed points for ρ_t^i, W_{t+1}^i, N_t^i, ρ_t^{-i}, W_{t+1}^{-i}, and N_t^{-i}.

3.1.1 Numerical example: rejectable offers

The discussion thus far has outlined a single backwards induction step to update the values $R_t^i(s)$, $M_t^i(s)$, and N_t^i, and likewise for $-i$. Iterating backwards in time until all value functions have converged, we arrive at an equilibrium of the game.[19] Figure 8 illustrates a numerical example, for two risk-neutral agents that split a cake of size $K = 100$, under a parameterization chosen for computational convenience: $\kappa = 0.2$, $\delta = 0.0005$, and $\bar{\rho} = 0.005$, and $f(h) = Zh^\zeta$ with $Z = 0.01$ and $\zeta = 2/3$. A discrete, evenly-spaced grid Γ^s of offers is considered, from $s = 0\%$ to $s = 100\%$ by steps of size 0.1%. For simplicity, we assume a uniform default distribution on the grid Γ^s:

$$\vec{\eta} \;=\; \vec{u} \;\equiv\; (1/n, \ldots, 1/n) \tag{91}$$

where $n = 1001$ is the number of points in the grid. Likewise, we assume a uniform default distribution for the acceptance choice, setting $(1 - \bar{\alpha}, \bar{\alpha}) = (1/2, 1/2)$. The decision arrival rate $\bar{\rho}$ is set low enough so that the probability of completing a decision in a single time step is close to zero; together with the low discount rate, it reflects our assumption that time steps are very brief. The offer step size is chosen fine enough to guarantee that equilibrium is unique

[19] Working backwards requires an initial guess of the value function at some hypothetical future time T, A good guess under a symmetric parameterization is $N_T^A = N_t^B = 0.49K$, $R_T^A(s) = \max\{sK, N_T^A\}$, and $R_T^A(s) = \max\{(1 - s)K, N_T^A\}$. But as long as the step size in the grid Γ^s is sufficiently fine, backwards induction converges to the same equilibrium regardless of the initial guess.

(for further discussion, see Sec. 3.5 below). The precision cost κ is chosen large enough so that errors are visible in the graphs, without displaying a high degree of irrationality. The value of Z is a normalization, but its setting relative to cake size K matters: crucially, we wish to consider parameters such that negotiation is actually worthwhile. Under a symmetric parameterization, this requires that the present discounted value of never bargaining, and instead devoting all future time to alternative activities, $Z/\delta = 20$, 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 shown as functions of the share s promised to 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 $R^B(1-s) = R^A(s)$; likewise all the functions describing player B's behavior are left-to-right reflections of the corresponding objects shown in Fig. 8 for player A.[20] 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 rises approximately linearly to a peak at $M^A(51.2) = 50.40$. It then drops slightly to a flat plateau at $M^A(s) = 47.69$ for offers in the range $s \geq 55$. 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 last panel of the figure. Higher offers would be more valuable to A if B accepted them, but the acceptance probability $\pi^B(s)$ is less than one percent for $s \geq 54$, so for large s, $M^A(s)$ mostly reflects the value of waiting for someone to formulate a new offer after the outstanding offer is rejected.

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 can then accept, ending the game. Therefore, the green curve $R^A(s)$, representing A's value of responding to an offer from B, lies just below the 45^o line for $s \geq 50$. And when $s \leq 46$, we have $R^A(s) = 47.41$, representing the value of the option to reject. While it might appear on a first glance that $R^A(s) \geq M^A(s)$ for all s, this is not true around $s \approx 50$, where $R^A(50) = 49.10 < M^A(50) = 49.39$. 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 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. Hence the value of receiving an offer of a perfectly even split is slightly less than the value of having already made the same offer. Both of these values also exceed the initial value $N^A = N^B = 48.14$ (red star), reflecting the fact that if no offer is outstanding then players will need to think, and time will elapse, before the terminal payoff is received.

Thus we see that equilibrium outcomes are simple and highly rational, resembling the Rubinstein (1982) game; the actual split is typically near to 50/50, with a small advantage to the proposer, in spite of the presence of errors in the model. The modal proposal by A is 51.2, and 90% of all offers made by A lie in the interval [50,52] (top, right panel of Fig. 9). An even split is accepted with 99.6% probability, while A accepts the proposal $s = 49$ with 93.6% probability and $s = 48$ with 34.2% probability (last panel of the figure). However, some low-cost errors are present. For example, although it is hard to see in the top, right panel, A has a tiny, roughly constant probability (0.012%) of making any offer s in the range $s > 55$. These unacceptable offers are not very costly, because they simply postpone arrival of an agreement. On the other hand, the probability that A makes an offer below 50 declines extremely rapidly; only one in a

[20] This equilibrium was computed by initializing with a symmetric guess of the form $N_T^A = N_T^B$. But as long as the step size in the grid Γ^s is sufficiently small, the same (unique) equilibrium is obtained starting from an asymmetric guess. See Sec. 3.5.

Figure 8: Equilibrium of rejectable offers game.

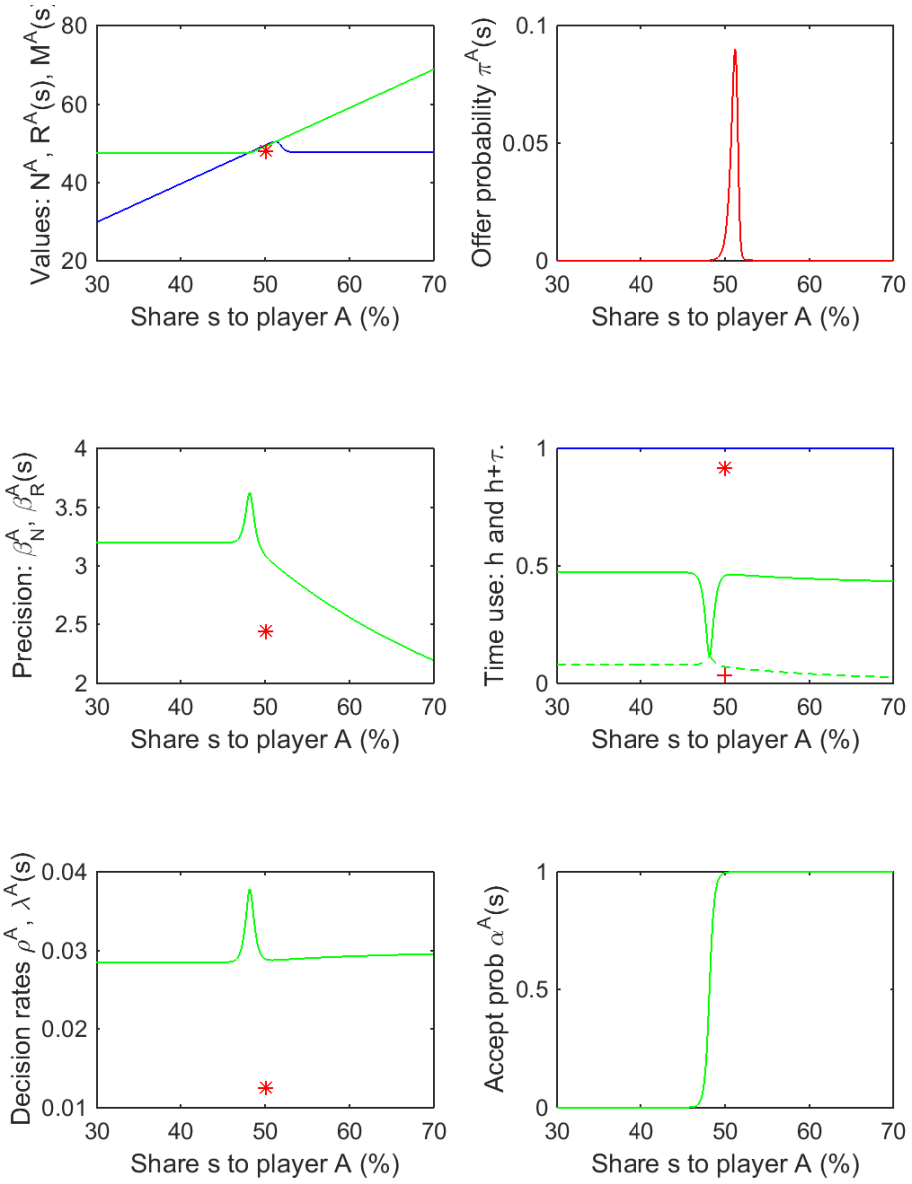

Note: Player A behavior in equilibrium of rejectable offers game (choices described by Prop. 4).
Objects at node N in red; node R in green; node M in blue.

Top left: Values to A of outstanding offers. Red star: Value $N^A = 48.18$ when no offer has been made. Green: Value $R^A(s)$ of offer made by B. Blue: Value $M^A(s)$ of offer made by A.
Top right: Probabilities $\pi^A(s)$ of offers s made by A at node N^A.
Middle left: Precision of A. Red: offer precision β_N^A at node N^A. Green: response precision $\beta_R^A(s)$.
Middle right: Time use of A at nodes $R^A(s)$ and N^A. Dash and cross: labor h; solid and star: $h + \tau$.
Bottom left: Decision arrival probabilities of player A: ρ^A (red) and $\lambda^A(s)$ (green).
Bottom right: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

million offers made by A lies in the range $s \leq 45.5$. Low offers are very costly mistakes to A, since B accepts them with very high probability.

While behavior in this example is close that found under full rationality, the bottom panels of Figure 9 show the obvious contrast with Rubinstein (1982): agreement is not immediate. Instead, offers take time to arrive, and their order of arrival is an endogenous result of the model. The bottom right panel shows that, conditional on responding, A accepts any offer $s \geq 50$ with probability α close to one. However, the bottom left panel shows that this response takes time, arriving with probability $\lambda \approx 3\%$ in any single time step. This decision rate corresponds to spending roughly 40% of time choosing whether to accept or reject the outstanding offer (the vertical distance between the solid and dashed green lines in the second panel of the second row shows that $\tau \approx 0.4$). This time allocation remains roughly constant across offers s, except near $s = 50$, where the player is nearly indifferent between accepting and rejecting, so τ falls sharply, and instead the player devotes effort to finishing the decision quickly (we see a spike in the vertical distance μ above the solid green line, leading to the spike in the arrival rate ρ shown in the bottom left panel). In contrast, each player's offer arrival rate ρ is roughly 1% per time step at node N (red star in the bottom left panel). This slower arrival rate reflects the fact that making an offer requires the player to consider many options, rather than just two. Hence at node N, almost all time is devoted to choosing between possible offers ($\tau \approx 0.9$ at node N, seen as the vertical distance between the red star and the red cross in the second row, second panel).

Thus, taking into account both player's decisions, around fifty time steps are expected to pass before the first offer is made, and another thirty elapse before it is (usually) accepted. Indeed, the game usually ends with the acceptance of the first offer: the probability of this event is $\sum_{s \in \Gamma^s} \pi^i(s)\alpha^{-i}(s) = 82.5\%$. This is reminiscent of the fact that the first offer is accepted immediately in Rubinstein's equilibrium; but here there is a delay, and a non-negligible probability that the first offer fails.

3.2 Solving the alternating offers game

The game we just solved assumed that the response to any offer took the form "yes" or "no". We next consider the alternating offers protocol, in which rejecting an offer requires a player to make a counteroffer. The game tree at node N_t^i is still the one seen in Fig. 5. The subgames at nodes $M_t^i(s)$ and $R_t^i(s)$ are changed by eliminating the rejection option that leads back to N_{t+1}^i, and substituting it with a choice across many counteroffers $s' \in \Gamma^s$.

Defining notation for the equilibrium policies at node $R_t^i(s)$ of this game, let $h_t^i(s)$ be time devoted to "work", let $\lambda_t^i(s)$ be the response arrival rate, let $\alpha_t^i(s)$ be the probability of accepting an offer, and let $\gamma_t^i(s'|s)$ be the probability of a given counteroffer s'. Both $\alpha_t^i(s)$ and $\gamma_t^i(s'|s)$ are defined conditional on arrival of a response, so that

$$\alpha_t^i(s) + \sum_{s' \in \Gamma^s} \gamma_t^i(s'|s) \;=\; 1. \tag{92}$$

Then the value of responding to an offer satisfies:[21]

$$R_t^i(s) \;=\; f(h_t^i(s)) + (1-\delta)\left\{\lambda_t^i(s)\left(\sum_{s'\in\Gamma^s}\gamma_t^i(s'|s)M_{t+1}^i(s') + \alpha_t^i(s)u^i(s)\right) + (1-\lambda_t^i(s))R_{t+1}^i\right\}, \tag{93}$$

and the value waiting for a response to one's own outstanding offer is

$$M_t^i(s) \;=\; f(1) + (1-\delta)\left\{\lambda_t^{-i}(s)\left(\sum_{s'\in\Gamma^s}\gamma_t^{-i}(s'|s)R_{t+1}^i(s') + \alpha_t^{-i}(s)u^i(s)\right) + (1-\lambda_t^{-i}(s))M_{t+1}^i\right\}. \tag{94}$$

A numerical example of alternating offers equilibrium is shown in Fig. 9. Parameters are identical to the previous example, except that we must define a default distribution compatible with the choice set at node $R_t^i(s)$ under alternating offers. To make the results quantitatively comparable with those from the rejectable offers case, we hold fixed the default probability $\bar{\alpha}$ associated with the option to accept. And conditional on making a counteroffer, we impose the same default offer probabilities that apply in state N_t^i. In other words, we benchmark the choice of offering any $s\prime \in \Gamma^s$, or accepting the outstanding offer s, against the following default probabilities:

$$((1-\bar{\alpha})\vec{u}, \bar{\alpha}) \;=\; \left(\frac{1-\bar{\alpha}}{n}, \ldots, \frac{1-\bar{\alpha}}{n}, \bar{\alpha}\right), \tag{95}$$

with $\bar{\alpha} = 1/2$ and $n = 1001$ as before.

The value functions, the initial offer probabilities, and the acceptance probabilities are very similar to those we saw before. As in the previous example, the offer that is actually accepted remains close to a 50/50 split, with a small advantage to the proposer. A's modal offer is now $s = 50.8$, and the probability that A's offer lies in [50,52] is 92%. As before, the game is likely to end with the acceptance of the first offer; the probability is 83.8%, slightly higher than in the rejectable offers case. Also, note that the second panel of the top row shows A's counteroffer probabilities $\gamma^A(s'|s)$ alongside the initial offer probabilities $\pi_t^A(s)$ associated with node N_t^A. The counteroffer probabilities (green, graphed conditional on an outstanding offer $s = 0.5$) are so similar to the initial offer probabilities that green curve is almost entirely obscured by the red curve representing $\pi_t^A(s)$. Thus, the distribution of offers remains similar even after a first offer is rejected.

The most visible differences in this new specification relate to decision rates and time use. The first offer arrives at the rate $\rho_t^A \approx 1\%$, and responses to "good" offers arrive with roughly 3% probability per time step, as in the rejectable offers case. But rejecting a "bad" offer is now a harder decision, since it requires a counterproposal. Therefore the response rate to low offers drops back to $\lambda_t^A(s) \approx 1\%$, similar to the speed of making the first offer. Similarly, looking at time use (second row, second panel), we see that roughly 90% of A's time is dedicated to formulating an offer in state N_t^A (vertical distance between red star and red cross) and when A is facing a low offer (vertical distance between solid and dashed green lines). But when facing a good offer, A reallocates time from choosing a response (τ falls to 0.4), to concentrating on finishing the decision quickly (μ rises to roughly 0.5, and therefore the decision rate rises to 3% per time step).

[21] Here, for brevity, we just state the value function in terms of the policy functions, instead of spelling out the underlying maximization problem.

Figure 9: Equilibrium of alternating offers game.

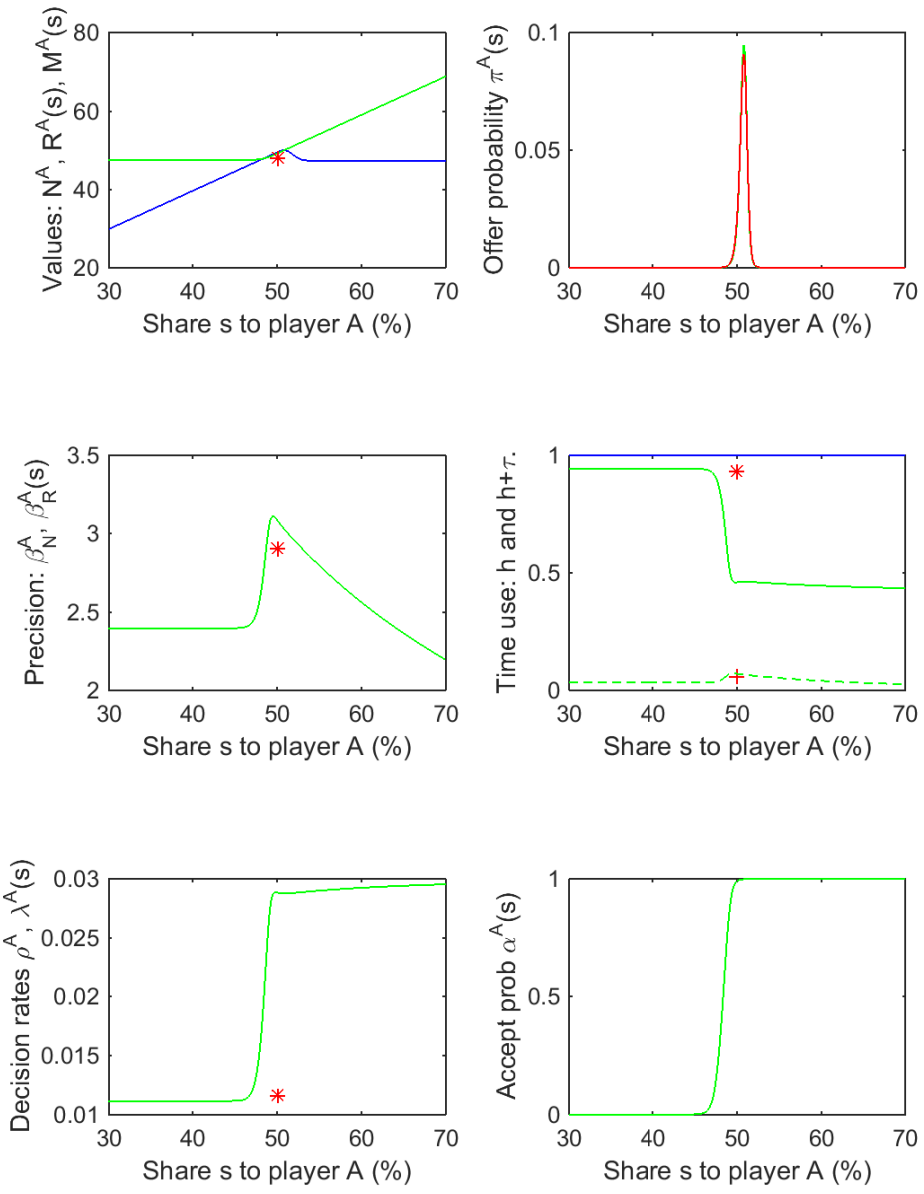

Note: Player A behavior in equilibrium of alternating offers game (choices described by Prop. 4).
Objects at node N in red; node R in green; node M in blue.

Top left: Values to A of outstanding offers. Red star: Value $N^A = 48.16$ when no offer has been made. Green: Value $R^A(s)$ of offer made by B. Blue: Value $M^A(s)$ of offer made by A.
Top right: Probabilities $\pi^A(s)$ and $\gamma^A(s'|s)$ of offers s at node N^A (red) and counteroffers s' at $R^A(s)$ (green).
Middle left: Precision of A. Red: offer precision β^A_N at node N^A. Green: response precision $\beta^A_R(s)$.
Middle right: Time use of A at nodes $R^A(s)$ and N^A. Dash and cross: labor h; solid and star: $h + \tau$.
Bottom left: Decision arrival probabilities of player A: ρ^A (red) and $\lambda^A(s)$ (green).
Bottom right: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

3.3 Withdrawable or updateable offers

Withdrawable offers protocol

The sequence of choices in the withdrawable bargaining protocol is the same as in the rejectable offers protocol (illustrated by Figs. 5-7), except that agent i at node $M_t^i(s)$ has the option of withdrawing the outstanding offer, thus returning the game to node N_t^i, where no offer is outstanding. Note that choosing to withdraw does not immediately entail a choice across any other alternatives. Therefore player i's withdrawal decision at $M_t^i(s)$ is an example of an error-prone stopping problem — as analyzed in Sec. 2.2 — instead of the error-prone decision problem analyzed in Sec. 2.3.

For comparability with previous results, the default withdrawal rate is set equal to the default decision rate $\bar{\rho} = 0.005$ assumed in previous simulations. All other default probabilities take the same values that were assumed in the rejectable offers case. We write the equilibrium withdrawal probability as $\xi_t^i(s)$, with notation for the other policy functions as before. Then the value function at node $M_t^i(s)$ satisfies

$$\begin{aligned} M_t^i(s) \;&= f(h_t^i(s)) + (1-\delta)\xi_t^i(s)N_{t+1}^i \\ &\quad + (1-\delta)(1-\xi_t^i(s))\left\{\lambda_t^{-i}(s)\left((1-\alpha_t^{-i}(s))N_{t+1}^i + \alpha_t^{-i}(s)u^i(s)\right) + (1-\lambda_t^{-i}(s))M_{t+1}^i(s)\right\}, \end{aligned} \tag{96}$$

Likewise, the option to withdraw alters the equation for the value function $R_t^i(s)$:

$$\begin{aligned} R_t^i(s) \;&= f(h_t^i(s)) + (1-\delta)\lambda_t^i(s)\left((1-\alpha_t^i(s))N_{t+1}^i + \alpha_t^i(s)u^i(s)\right) \\ &\quad + (1-\delta)(1-\lambda_t^i(s))\left((1-\xi_t^{-i}(s))R_{t+1}^i(s) + \xi_t^{-i}(s)N_{t+1}^i\right). \end{aligned} \tag{97}$$

To evaluate equations (96)-(97), i's decision variables $h_t^i(s)$, $\lambda_t^i(s)$, $\gamma_t^i(s_j|s)$, $\alpha_t^i(s)$, and $\xi_t^i(s)$, as well as the corresponding policies of $-i$, are calculated by applying Prop. 4.[22]

Updateable offers protocol

Like the withdrawable offers protocol, the updateable offers protocol assumes it is possible to take back an offer that has already been made. But in this case the player that made the existing offer can only cancel it by proposing an alternative. Thus, it is a natural extension of the alternating offers protocol. We assume the same default probabilities that we considered in the alternating case, extending them in the most obvious way: the default decision rate and default offer distribution for updating one's own offer are set equal to those associated with a making counteroffer.

Thus, let $\xi_t^i(s)$ be the probability of updating one's own current offer s with a new offer s', and let $\upsilon_t^i(s'|s)$ be the distribution of updated offers. Then equation (94) is replaced by the following equation for $M_t^i(s)$:

$$\begin{aligned} M_t^i(s) \;&= f(h_t^i(s)) + (1-\delta)\xi_t^i(s)\sum_{s'\in\Gamma^s}\upsilon_t^i(s'|s)M_{t+1}^i(s') \\ &\quad + (1-\delta)(1-\xi_t^i(s))\left\{\lambda_t^{-i}(s)\left(\sum_{s'\in\Gamma^s}\gamma_t^{-i}(s'|s)R_{t+1}^i(s') + \alpha_t^{-i}(s)u^i(s)\right) + (1-\lambda_t^{-i}(s))M_{t+1}^i(s)\right\}, \end{aligned} \tag{98}$$

[22] Prop. 2 is a special case of Prop. 4. Therefore the program `twomargins_eta.m` that solves the general model of Sec. 2.3 is also applicable to the special case of an error-prone stopping problem.

Figure 10: Equilibrium of game with withdrawable offers.

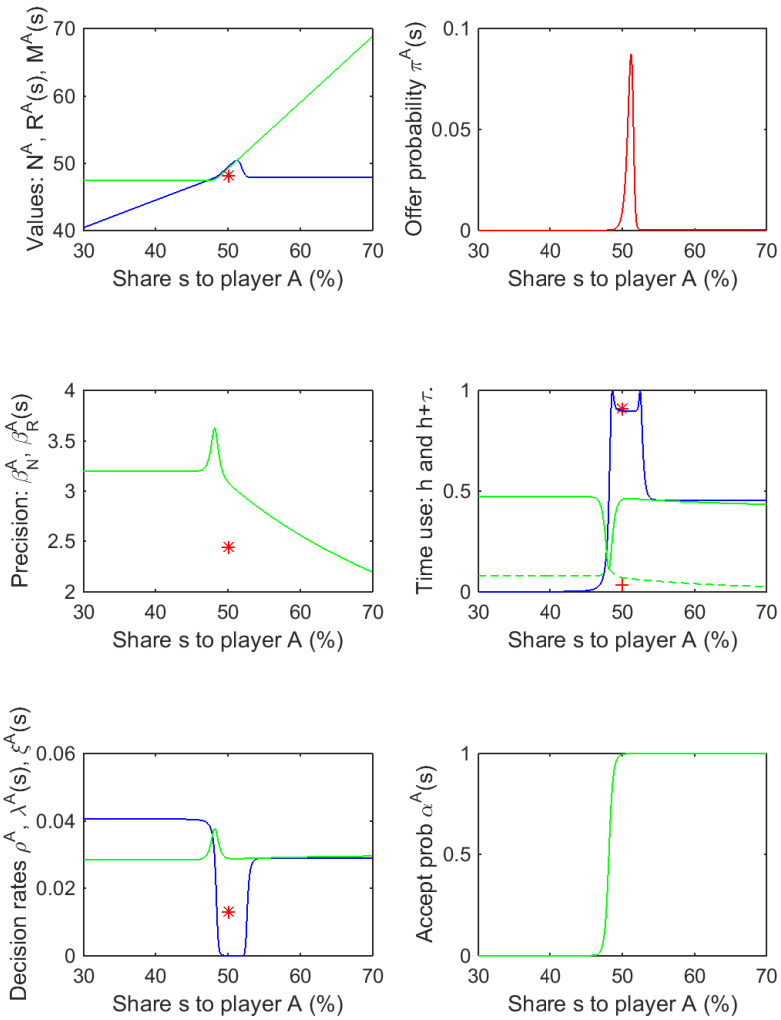

Note: Player A behavior in equilibrium of withdrawable offers game (choices described by Prop. 4). Objects at node N in red; node R in green; node M in blue.

Top left: Values to A of outstanding offers. Red star: Value $N^A = 48.17$ when no offer has been made. Green: Value $R^A(s)$ of offer made by B. Blue: Value $M^A(s)$ of offer made by A.
Top right: Probabilities $\pi^A(s)$ of offers s made by A at node N^A.
Middle left: Precision of A. Red: offer precision β_N^A at node N^A. Green: response precision $\beta_R^A(s)$.
Middle right: Time use of A at nodes $R^A(s)$ and N^A. Dash and cross: labor h; solid and star: $h + \tau$.
Bottom left: Decision arrival probabilities of player A: ρ^A (red) and $\lambda^A(s)$ (green).
Bottom right: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

The notation for the other policy functions is as before. Likewise, the fact that $-i$ can update the outstanding offer is reflected in the equation for $R_t^i(s)$, which becomes

$$R_t^i(s) = f(h_t^i(s)) + (1-\delta)\lambda_t^i(s)\left(\sum_{s'\in\Gamma^s}\gamma_t^i(s'|s)M_{t+1}^i(s') + \alpha_t^i(s)u^i(s)\right)$$
$$+ (1-\delta)(1-\lambda_t^i(s))\left((1-\xi_t^{-i}(s))R_{t+1}^i(s) + \xi_t^{-i}(s)\sum_{s'\in\Gamma^s}\upsilon_t^{-i}(s'|s)R_{t+1}^i(s')\right). \quad (99)$$

Results

The results of the withdrawable and updateable protocols are shown in Figures 10 and 11, respectively. These graphs show a few more curves than Figs. 8-9, corresponding to the additional decisions that play a role in the updateable and withdrawable specifications. Nonetheless, a brief inspection shows close similarities with our previous results. In particular, the curves shown in Fig. 8 for the rejectable offers protocol are almost unchanged in Fig. 10, which illustrates the withdrawable offers protocol. Likewise, the curves seen in Fig. 9 (alternating offers) are almost identical to the corresponding curves in Fig. 11 (updateable offers). The most noteworthy difference is seen in the value function $M^A(s)$, seen in blue in the top left panels of Figs. 8-11. If offers can be withdrawn or updated, then an excessively low offer is a less costly mistake. Therefore, the curve $M^A(s)$ is higher and less steeply sloped over the range $s < 50$ in Figs. 10-11 than it was in Figs. 8-9, which described protocols under which rescinding an offer was impossible.

The additional curves seen in Fig. 10 (in blue) illustrate the withdrawal decision at node $M^A(s)$. The blue curve in the middle right panel illustrates time use at $M^A(s)$: the height of the curve represents "work", and the vertical distance above the curve represents time devoted to the withdrawal decision. When player A has made a reasonable offer that is likely to be accepted ($s \approx 50$), she devotes almost all her time to work ($h \approx 0.9$). But when the outstanding offer is slightly below 50, her time use shifts dramatically, concentrating entirely on the withdrawal decision ($h \approx 0$); thus the probability of withdrawal rises from near zero to roughly 4% per period (blue curve in bottom left panel). When the outstanding offer is instead above 50, A devotes roughly half her time to the withdrawal decision ($h \approx 0.5$); an excessively high offer is a less costly error than an excessively low one, but player A still stands to gain by rescinding her unacceptable offer, so she withdraws it with roughly 3% probability per period.

The interpretation of the blue curves in Fig. 11 (updateable offers protocol) is similar. In this protocol there are three uses of time at node $M^A(s)$, as seen in the middle right panel: the height of the dashed blue line represents h ("work"), the distance above the solid blue line represents μ (time devoted to choosing whether to update the current offer), and the distance between the solid and dashed lines represents τ (time devoted to deciding which new offer to make). Again we see that when the outstanding offer is sufficiently close to 50, almost all time is spent working. But for $s < 50$, h falls rapidly to zero as player A devotes time to decision-making instead, and the arrival rate of an updated offer is between 3% and 4% per period. For $s > 50$, time use is divided more or less equally between work and decision-making, and an updated offer arrives with roughly 1.5% probability per time step.

The top right panel of Fig. 11 also features a blue curve which represents the distribution $\upsilon^A(s'|s)$ of updated offers s'. We illustrate this distribution conditional on the outstanding offer $s = 0.45$, a relatively bad offer which player B is eager to accept and player A is eager to update.

Figure 11: Equilibrium of game with updateable offers.

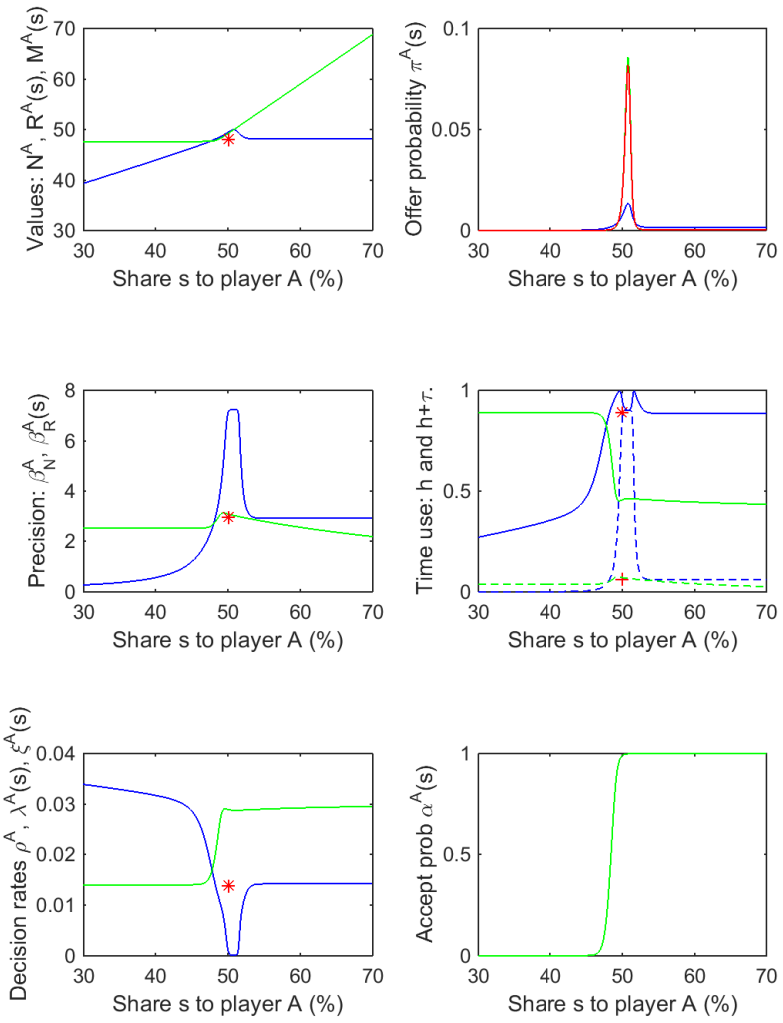

Note: Player A behavior in equilibrium of updateable offers game (choices described by Prop. 4).
Objects at node N in red; node R in green; node M in blue.

Top left: Values to A of outstanding offers. Red star: Value $N^A = 48.15$ when no offer has been made. Green: Value $R^A(s)$ of offer made by B. Blue: Value $M^A(s)$ of offer made by A.
Top right: Probabilities $\pi^A(s)$, $\gamma^A(s'|s)$, and $\xi^A(s'|s)$ of offers s at node N^A (red) and counteroffers s' at $R^A(s)$ (green) and $M^A(s)$ (blue).
Middle left: Precision of A. Red: offer precision β_N^A at node N^A. Green: response precision $\beta_R^A(s)$.
Middle right: Time use of A at nodes $R^A(s)$ and N^A. Dash and cross: labor h; solid and star: $h + \tau$.
Bottom left: Decision arrival probabilities of player A: ρ^A (red) and $\lambda^A(s)$ (green).
Bottom right: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

As the middle right panel shows, conditional on $s = 0.45$ player A devotes most of her time to deciding to rescind the offer, and less time to choosing an appropriate counteroffer ($\mu >> \tau$ at $s = 0.45$). That is, in her hurry to correct a poor offer, player A hastily proposes an alternative. Therefore the precision of the updated offer is low, which explains why the blue curve in the top right panel (the distribution υ^A of updated offers s') is much less sharply peaked than the red curve (the distribution π of initial offers s). This can also be seen in the middle left panel, which shows the precision β of the distribution of counteroffers (in blue). Updated offers are most precise when the outstanding offer s is close to the optimum (just above 50), and rapidly become imprecise otherwise, especially if the outstanding offer is too low instead of too high.

While the strategies played under the withdrawable and updateable protocols are more complex than those associated with the first two frameworks, ultimately the most striking conclusion about the cases compared is the similarity of the equilibrium distribution of outcomes in all four protocols. The first offer is expected to arrive after approximately 50 time steps. Initial offers are tightly clustered around a 50/50 split in all specifications, with a small advantage to the proposer. Offers in this range are rarely withdrawn or updated; instead they are typically accepted after roughly 30 periods. The probability that the first offer is accepted, thus ending the game, is 80.3% and 72.6% under the withdrawable and updateable protocols, respectively, somewhat lower than the probabilities for the other protocols considered. Stationarity of the game implies that under the rejectable and withdrawable protocols, subsequent offers are drawn from the same distribution as the first if the first is rescinded. The distributions of counteroffers in the alternating and updateable protocols, and the arrival rates of these counteroffers, are also very similar to the distribution of initial offers. When an update occurs under the updateable protocol, it is drawn slowly from a more tightly clustered distribution if the outstanding offer was already close to the optimal offer, and it is drawn quickly with lower precision if the outstanding offer was excessively generous; but updates are unlikely events *ex ante*.

The unimportance of changes in the bargaining protocol recalls and reinforces some intuitions from full-information games. Indeed, Perry and Reny (1993) argue that under full rationality, bargaining protocols with rejection and with proposal of alternatives are equivalent; therefore they focus only on the latter. While these protocols are not formally equivalent in the presence of costly, error-prone choice, nonetheless our results show that they are very similar when costs are parameterized as we have done here. Likewise, the ability to rescind an offer (prior to acceptance) plays no role in a fully rational game, since no player will make an offer that he perceives to be suboptimal. In a bargaining game where players may make mistakes, the ability to rescind an offer may obviously be valuable *ex post* when an error occurs. Nonetheless, our results suggest that the possibility of rescinding offers also has little effect on the value of a game *ex ante*, as long as players have the ability to exert cognitive effort in order to avoid costly errors.

3.4 Comparative statics

We next explore how our results change as parameters vary. Figure 12 illustrates a comparative statics exercise in which we make time inherently more valuable, raising the productivity of alternative activities from its benchmark value of $Z = 1$ to $Z = 2$. (See the Appendix for the graphs in this subsection and the following one.) The setup and the parameters are otherwise identical to the rejectable offers benchmark of Fig. 8. Since time is more costly, less of it is devoted to the decision. Comparing the red cross in the middle-right panel of Fig. 12 to that

of Fig. 8, the time devoted to alternative activities rises from only $h = 0.035$ when $Z = 1$ to $h = 0.397$ when $Z = 2$. Therefore the decision arrives more slowly, and with less accuracy. In particular, the probability that Z offers $s \in [49, 52]$ falls from 93.4% in the benchmark specification to 86.6% when $Z = 2$. 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 $Z = 2$ (red star in bottom left panel of the figure), 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 $Z = 1$, 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 0.005 to 0.001 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 0.01 raises the arrival rate of initial offers to 1.5% per time step.

A rather different comparative statics exercise is considered in Fig. 13, of alternative activities is raised to $Z = 3$. The interesting aspect 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 this "rotten pie". That is, the value of "working" full time forever, $h(1)/\delta = Z/\delta = 60$ exceeds the value of half the pie (50) in this example.

The results change in a sensible way. If these players were fully rational, they would simply avoid negotiating altogether. Being imperfectly rational, they are not entirely sure whether they should think about negotiations. Therefore they do spend some time analyzing the bargaining game, but the red dot in the middle-right panel of Fig. 13 shows that almost all their time is devoted to "work" ($h = 0.95$). Therefore offers arrive very slowly; initial offers arrive with probability 0.45% per time step (red star in the bottom-left panel). Since ending the game by accepting a 50/50 share is worse than continuing without agreement, the probability of offering or 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 accept offers above $s = 40$, on the rare occasions when A does make offers, some of these are as generous as $s = 42.5$. Indeed, the two panels in the top row show that A is almost indifferent about what offer she makes as long as her offer will be rejected almost certainly by B.

3.5 When is equilibrium unique?

The examples analyzed thus far, which assumed symmetric parameters for the two players, were calculated from a symmetric initial guess, and yielded symmetric equilibria. However, we have not yet investigated whether other equilibria might also exist.

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 proposes 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 symmetric parameters, the expected payoffs of A and B may differ.

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 single discrete time step, such a rapid choice implies low precision but requires a high fraction of time devoted to monitoring, so this is usually suboptimal. Knowing that i might complete an offer over the interval $[t, t+1]$ with a small positive probability does not suffice to rule out decision effort by player $-i$. Therefore, in equilibrium at node N^i, both players think simultaneously about making offers.

The possibility of multiplicity of expected shares is a separate issue from multiplicity in timing. To investigate this question, we now report some examples computed from asymmetric initial guesses, and check whether they converge to a symmetric equilibrium. Our baseline simulation computed the rejectable offers scenario from a symmetric initial guess, $N_T^A = N_T^B = 50$, which converged to the symmetric equilibrium shown in Sec. 3.1, with $N^A = N^B = 48.1811$. To check for multiplicity, we have run the same calculation from a highly asymmetric initial guess, $N_T^A = 80$, $N_T^B = 20$. Under the baseline parameterization — which assumes a fine grid Γ^s on $[0, 100]$ by steps of size 0.1 — we find that the asymmetry in players' expected shares unravels as backwards induction proceeds, converging finally to the same symmetric equilibrium, with $N^A = N^B = 48.1811$. The left panel of Figure 14 shows the value functions from both simulations. The baseline simulation that starts from a symmetric guess is plotted with $M^A(s)$ in blue and $R^A(s)$ in green. The simulation that starts from an asymmetric guess is plotted showing 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 high share (represented by a high value of N_{t+1}^A), she will reject low offers, so N_t^A cannot lie far below the assumed value of N_{t+1}^A. 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 right panel of Figure 14, which shows the value functions that are calculated from the asymmetric initial guess $N_T^A = 80$, $N_T^B = 20$, conditional on four different offer grids Γ^s. 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 an asymmetric equilibrium with $N^A = 53.10$, $N^B = 43.57$. Given a step size of 1.0, the computation instead converges to $N^A = 50.03$, $N^B = 46.37$, shown with dashed lines in the graph. Finally, graph also shows equilibria calculated with a step size of 0.5 (dash-dot lines) and with step size 0.1 (the benchmark parameterization, shown with solid lines). Both of these equilibria converge to full symmetry, with $N^A = N^B = 48.1811$, even when we start from an asymmetric guess.[23] While the equilibria with step sizes 0.5 and 0.1 are both plotted, they are indistinguishable in the graph.

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, backwards induction may instead 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

[23] The solid curves are identical to those shown in the left panel of the figure.

the game is defined on a continuous space of possible offers, but demonstrating this rigorously is beyond the scope of the numerical analysis of this paper.

Additional exercises were also performed (graphs available on request) that lowered the decision noise parameter from its benchmark value of $\kappa = 0.2$ to $\kappa = 0.02$, in order to explore how the degree of rationality affects multiplicity of equilibrium. We find that if decisions are less costly, then a wider range of equilibrium shares can be sustained, conditional on a given discrete grid. Assuming $\kappa = 0.02$ and a step size of 2.5, backwards induction from a guess favorable to A converges to $N^A = 76.7$, $N^B = 22.4$. With step size 0.5, the same calculation converges to $N^A = 62.1$, $N^B = 37.3$; and even with step size 0.25, it converges to $N^A = 54.8$, $N^B = 44.6$. We only find uniqueness when we go all the way to step size 0.1; then backwards induction from an asymmetric starting guess converges all the way to a symmetric equilibrium with $N^A = N^B = 49.66$.

4 Conclusions

This paper has presented a model of a time-consuming choice across several feasible actions, in which the decision-maker may choose quickly, achieving low precision, or more slowly, achieving arbitrarily high precision. Taking into account the time cost of choice, the decision-maker prefers to make some "errors"; it is not optimal to achieve perfect precision by picking the option with the highest gross payoff with probability one.

Besides errors in the choice across feasible actions, the paper also considers errors in the timing of decisions. Several cases 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 decision-making is embedded inside a game, the time devoted to choice must be reflected in the extensive form. Error-prone choice is known to be helpful for predicting play in extensive form games (*e.g.* McKelvey and Palfrey, 1998; Anderson, Goeree, and Holt, 2002). But while previous work focused on games in which the timing of decisions was fixed exogenously at the outset by the extensive form, in our setup costly, error-prone choice endogenizes the timing of decisions. This could prove to be a useful extension of the control cost equilibrium concept, and it would be interesting to test its performance in the laboratory, especially in experiments where players make intermittent decisions in continuous time.

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. Several other protocols are also considered, in which rejecting an offer requires the formulation of an alternative offer, or in which the proposer may withdraw or update the offer.

Numerical simulations of these bargaining games reveal an equilibrium structure very similar to that of Rubinstein (1982), except that agreement is not achieved instantaneously, and offers are sometimes rejected. Under a symmetric parameterization, there is a high probability of a roughly 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 slightly favorable to

themselves, and are usually willing to accept shares that are 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. While their game exhibited two types of multiplicity of equilibrium, we instead find a unique equilibrium as long as offers are drawn from a sufficiently fine grid. Assuming higher noise in decision-making also makes uniqueness more likely. Thus we find that the results of the Rubinstein (1982) and Binmore, Rubinstein, and Wolinsky (1986) games are more robust than Perry and Reny's generalization suggested.

References

Anderson, S., A. de Palma, and J.-F. c. Thisse (1992): *Discrete Choice Theory of Product Differentiation*. The MIT Press.

Anderson, S., J. Goeree, and C. Holt (2002): "The logit equilibrium: a perspective on intuitive behavioral anomalies," *Southern Economic Journal*, 69(1), 21–47.

Baron, R., J. Durieu, H. Haller, and P. Solal (2002): "Control costs and potential functions for spatial games," *International Journal of Game Theory*, 31, 541–61.

Barro, R. (1977): "Long-term contracting, sticky prices, and monetary policy," *Journal of Monetary Economics*, 3, 305–16.

Binmore, K., A. Rubinstein, and A. Wolinsky (1986): "The Nash bargaining solution in economic modelling," *Rand Journal of Economics*, 17(2), 176–88.

Bono, J., and D. Wolpert (2009): "Statistical prediction of the outcome of a noncooperative game," American University Economics Working Papers 2009-20, American University.

——— (2010): "A theory of unstructured bargaining using distribution-valued solution concepts," American University Economics Working Papers 2010-14, American University.

Cheremukhin, A., P. Restrepo-Echevarria, and A. Tutino (2012): "The assignment of workers to jobs with endogenous information selection," Discussion paper, Federal Reserve Bank of Dallas.

Christiano, L., M. Eichenbaum, and C. Evans (2005): "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of Political Economy*, 113(1), 1–45.

Costain, J., and A. Nakov (2015): "Logit price dynamics," Discussion Papers 10731, CEPR.

Cover, T., and J. Thomas (2006): *Elements of Information Theory*. Wiley Interscience, 2 edn.

Gertler, M., L. Sala, and A. Trigari (2008): "An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining," *Journal of Money, Credit and Banking*, 40(8), 1713–64.

Goeree, J., and C. Holt (1999): "Stochastic game theory: for playing games, not just for doing theory," *Proc. Nat. Acad. Sci. USA*, 96, 10564–7.

——— (2001): "Ten little treasures of game theory and ten intuitive contradictions," *American Economic Review*, 91(5), 1402–22.

Hall, R., and P. Milgrom (2008): "The limited influence of unemployment on the wage bargain," *American Economic Review*, 98(4), 1653–74.

Hansen, L. P., and T. Sargent (2007): *Robustness*. Princeton Univ. Press.

Machina, M. (1985): "Stochastic choice functions generated from deterministic preferences over lotteries," *Economic Journal*, 95, 575–94.

Marsili, M. (1999): "On the multinomial logit model," *Physica A*, 269, 9–15.

Matejka, F., and A. McKay (2015): "Rational inattention to discrete choices: a new foundation for the multinomial logit model," *American Economic Review*, 105(1), 272–98.

Mattsson, L.-G., and J. Weibull (2002): "Probabilistic choice and procedurally bounded rationality," *Games and Economic Behavior*, 41, 61–78.

McKelvey, R., and T. Palfrey (1995): "Quantal response equilibrium for normal form games," *Games and Economic Behavior*, 10, 6–38.

——— (1998): "Quantal response equilibrium for extensive form games," *Experimental Economics*, 1, 9–41.

Merlo, A., and C. Wilson (1995): "A stochastic model of sequential bargaining with complete information," *Econometrica*, 63(2), 371–99.

——— (1998): "Efficient delays in a stochastic model of bargaining," *Economic Theory*, 11(1), 39–55.

——— (2010): "Identification of stochastic sequential bargaining models," Discussion paper, Univ. of Pennsylvania.

Moreno, D., and J. Wooders (1998): "An experimental study of communication and coordination in noncooperative games," *Games and Economic Behavior*, 24, 47–76.

Perry, M., and P. Reny (1993): "A non-cooperative bargaining model with strategically timed offers," *Journal of Economic Theory*, 59, 50–77.

Rubinstein, A. (1982): "Perfect equilibrium in a bargaining model," *Econometrica*, 50, 97–109.

——— (2013): "Response time and decision-making: an experimental study," *Judgment and Decision Making*, 8(5), 540–551.

Selten, R. (1975): "A reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4, 25–55.

Sims, C. (1998): "Stickiness," *Carnegie-Rochester Conference Series on Public Policy*, 49, 317–56.

——— (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50, 665–90.

Smets, F., and R. Wouters (2003): "An estimated stochastic dynamic general equilibrium model of the euro area," *Journal of the European Economic Association*, 1(5), 1123–1175.

Stahl, D. (1990): "Entropy control costs and entropic equilibrium," *International Journal of Game Theory*, 19, 129–38.

Theodorou, E., K. Dvijotham, and E. Todorov (2013): "From information-theoretic dualities to path integral and Kullback Leibler control: continuous and discrete time formulations," Discussion paper, 16th Yale Workshop on Learning and Adaptive Systems.

Todorov, E. (2009): "Efficient computation of optimal actions," *Proc. Nat. Acad. Sci. USA*, 106(28), 11487–11483.

Van Damme, E. (1991): *Stability and Perfection of Nash Equilibrium*. Springer Verlag, 2 edn.

Wolinsky, A. (1987): "Matching, search, and bargaining," *Journal of Economic Theory*, 42, 311–33.

Woodford, M. (2008): "Information-constrained state-dependent pricing," *Journal of Monetary Economics*, 56, S100–S124.

5 Appendix. Additional figures.

Figure 12: Equilibrium of rejectable offers game, assuming costlier time ($Z = 2$).

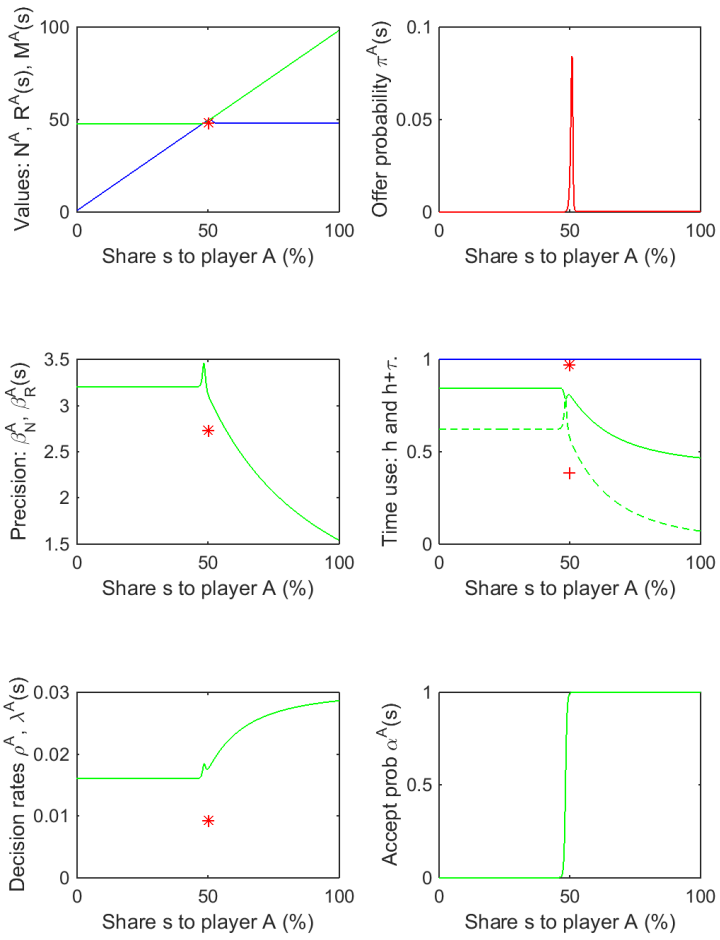

Note: Equilibrium of rejectable offers game (choices described by Prop. 4).
Green: Node R; blue: node M; red: node N. Assuming costlier time ($Z = 2$ instead of $Z = 1$).

Top left: Values to A of outstanding offers. Red: Value N^A when no offer has been made. Green: Value $R^A(s)$ of offer made by B. Blue: Value $M^A(s)$ of offer made by A.
Top right: Probabilities $\pi^A(s)$ of offers s made by A at node N^A.
Middle left: Precision of A. Red: offer precision β_N^A at node N^A. Green: response precision $\beta_R^A(s)$.
Middle right: Time use of A at nodes $R^A(s)$ and N^A. Dash and cross: labor h; solid and star: $h + \tau$.
Bottom left: Decision arrival probabilities of player A: ρ^A (red) and $\lambda^A(s)$ (green).
Bottom right: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

Figure 13: Equilibrium of rejectable offers game, with a "rotten pie" ($Z = 3$).

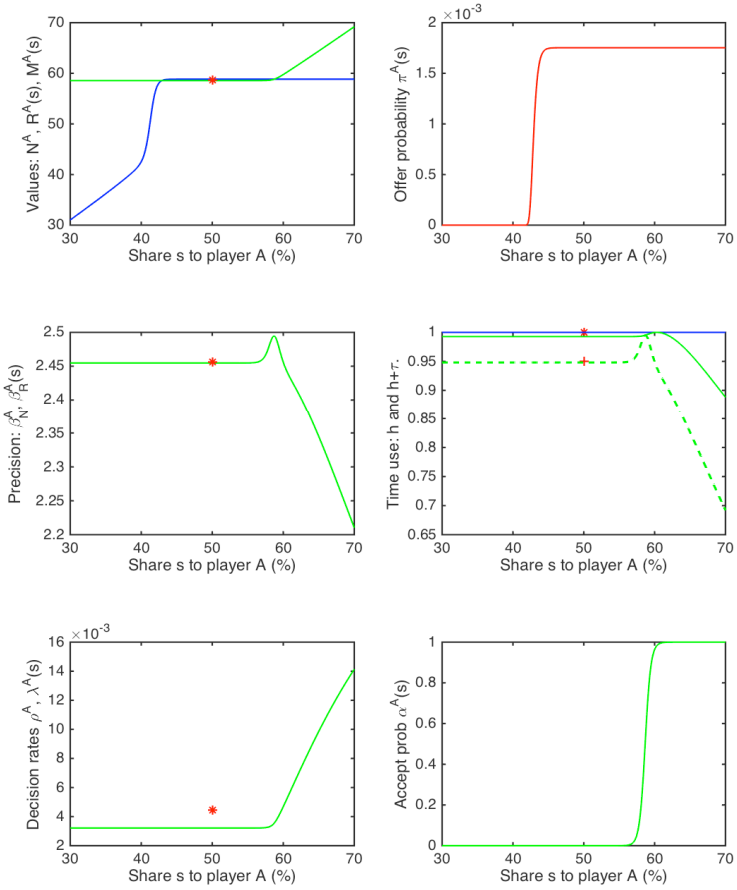

Note: Equilibrium of rejectable offers game (choices described by Prop. 4).
Green: Node R; blue: node M; red: node N. Time value raised to $Z = 3$, implying "rotten pie": $50 < Z/\delta$.

Top left: Values to A of outstanding offers. Red: Value N^A when no offer has been made. Green: Value $R^A(s)$ of offer made by B. Blue: Value $M^A(s)$ of offer made by A.
Top right: Probabilities $\pi^A(s)$ of offers s made by A at node N^A.
Middle left: Precision of A. Red: offer precision β^A_N at node N^A. Green: response precision $\beta^A_R(s)$.
Middle right: Time use of A at nodes $R^A(s)$ and N^A. Dash and cross: labor h; solid and star: $h + \tau$.
Bottom left: Decision arrival probabilities of player A: ρ^A (red) and $\lambda^A(s)$ (green).
Bottom right: Acceptance probabilities $\alpha^A(s)$ of player A, conditional on decision.

Figure 14: Equilibrium uniqueness and multiplicity in the rejectable offers game.

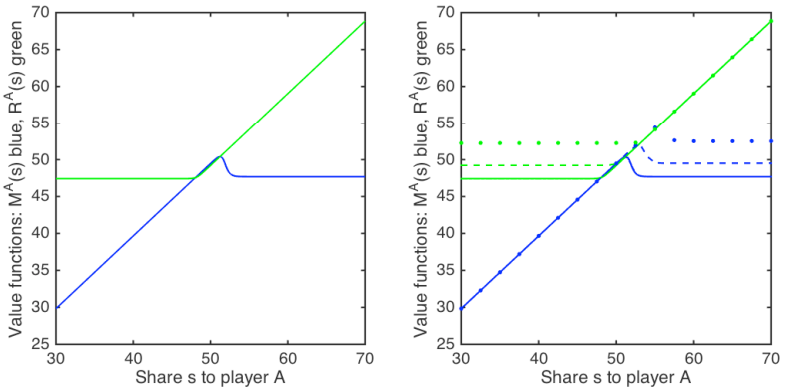

Note: Functions $R^A(s)$, in green, and $M^A(s)$, in blue, of rejectable offers game.

Left panel: Example of uniqueness. Two simulations shown.
Parameters same as benchmark equilibrium of Fig. 8, including step size 0.1 in grid Γ^s.
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.

Right panel: Examples of multiplicity when offer grid is coarse. Four simulations shown.
Parameters same as benchmark equilibrium of Fig. 8, except for step size in grid Γ^s.
Backwards induction from asymmetric guess with step size 0.1 (solid) or step size 0.5 (dash-dot) converges to symmetric equilibrium, indicating uniqueness.
Backwards induction from asymmetric guess with step size 1 (dash) or 2.5 (dot) converges to asymmetric equilibrium; indicates that multiple equilibrium are found, depending on initial guess.