Bayesian updating for self-assessment explains social dominance and winner-loser effects

Abstract

In animal contests, winners of previous contests often keep winning and losers keep losing. These so-called 'winner-loser effects', can have dramatic impacts on future contests and seem to play a key role in stabilizing the resulting dominance hierarchies. However, the cognitive mechanisms through which these effects occur are unknown. Winner-loser effects are often accompanied by a change in the aggressiveness of experienced individuals, which suggests individuals may be adjusting their self-assessment of their abilities based on this newly gained information. This updating of a prior estimate is ideally described by Bayesian updating. Here we show that Bayesian updating provides a mechanism to explain why winner-loser effects occur and makes clear predictions for the behavior of individuals and social groups. We implement an agent-based model to compare Bayesian updating to other possible strategies. We first show that Bayesian updating reproduces empirical results of the behavior of individuals and groups in dominance interactions. We then provide a series of testable predictions that can be used in future empirical work to distinguish Bayesian updating from other potential mechanisms. Our work demonstrates the utility of Bayesian updating as a mechanism to explain and ultimately, predict changes in behavior after salient social experiences.

Introduction

Nearly all social groups have measurable dominance hierarchies (Ellis 1995; Shizuka and McDonald 2015; Strauss et al. 2022). These hierarchies can be defined by consistent imbalances in pair-wise aggression and subordination (Rowell 1974; Drews 1993). In the simplest cases, dominance relationships are a measure of contest ability, with the stronger or otherwise more competitive individuals consistently winning over weaker ones. If all outcomes were determined entirely by the intrinsic ability to win a fight (Parker 1974) then dominance hierarchies would be perfectly linear and predictable. And while dominance hierarchies are typically very stable and more linear than would be predicted by chance (Jackson and Winnegrad 1988; Tibbetts et al. 2022), dominance outcomes, both in terms of individual contests and the overall positioning of particular individuals within the hierarchy, are often unpredictable (Landau 1951; Chase et al. 2002, 2022). This is because contest outcomes can be stochastic (reviewed in Strauss et al. 2022), and are influenced by a combination of factors, including motivation (Arnott and Elwood 2008; O'Connor et al. 2015), strategy (Parker 1974), individual state (Goetz et al. 2014; Li et al. 2014; Holekamp and Strauss 2016), and perhaps most importantly, social feedback and previous experiences (Chase et al. 2002).). This dichotomy—often unpredictable but unusually stable outcomes—highlights the gap between our current understanding of the mechanisms of dominance and the reality of natural systems. Despite the importance of social dominance for individual and group outcomes (Dewsbury 1988; Snyder-Mackler et al. 2020; Simons et al. 2022), we still lack models that can predict the formation of dominance hierarchies over time by incorporating the additional social

mechanisms that reinforce or even supersede differences in intrinsic contest ability (Chase et al. 2002; Tibbetts et al. 2022).

Winner and loser effects may be one of these social mechanisms governing the formation of dominance hierarchies (Dugatkin 1997; Bonabeau et al. 1999; Hemelrijk 2002). The winner effect is a widely observed phenomenon where success in one contest leads to increased probability of success in subsequent contests, while the loser effect describes how defeats often lead to more defeats (reviewed in Hsu et al. 2006). Although these two phenomena are tightly related, prior work shows that there is variation in the relative strength of winner and loser effects, and winner effects can even occur without loser effects (Goubault and Decuignière 2012), or vice versa (McDonald et al. 1968). Overall, the most typical phenomenon is for individuals to display both winner and loser effects, with the loser effect being relatively stronger (reviewed in Hsu et al. 2006; Rutte et al. 2006).

While in some contexts, winner and loser effects could be explained by changes in intrinsic, individual ability or opponent behavior, experiments tend to control for these effects such that changes in contest outcomes appear to be largely driven by a change in the focal individual's aggressive behavior (reviewed in Hsu, 2006). This change in aggression suggests that individuals may be modifying their self-assessment of their chances of winning in future contests based on the outcome of their current contest (Rutte et al. 2006). Understanding what motivates this change in behavior and how changes in self-assessment may occur is central to our understanding of winner-loser effects and dominance formation.

Most of our biological understanding of winner-loser effects relates to short-term endocrine-mediated shifts in behavior and the accompanying neurophysiology (reviewed in Zhou et al, 2018). For example, several studies in rodents showed that a boost in testosterone following a contest is sufficient to induce winner effects (Fuxjager et al. 2011*a*), however blocking testosterone merely attenuated the winner effect (Fuxjager et al. 2011*b*). Furthermore, winner effects are mediated by social and environmental context (Fuxjager and Marler 2010; Fuxjager et al. 2011*b*), independent of the effects of testosterone, while other studies have pointed to additional contributions from dopamine (Becker and Marler 2015; Ibuchi and Nagayama 2021) and serotonin (Winberg and Thörnqvist 2016; Hu et al. 2020).

Despite these putative mechanisms, our general understanding of the drivers of post-contest behavior remains limited, and even these mechanisms primarily explain short term winner-loser effects (but see Cordero and Sandi 2007; Winberg and Thörnqvist 2016). While winner-loser effects are often short-lived in experimental tests—generally lasting from mere minutes to days (Chase et al. 1994; Johnsson 1997) or until some external disruption occurs (Stevenson et al. 2000)—some studies have found persistent winner/loser effects influencing behavior up to months after the initial conflict (Lan and Hsu 2011; Laskowski et al. 2016; Momohara et al. 2016). A clear mechanistic model of the cognitive changes that occur following contests could help reconcile these seemingly disparate results and allow us to better predict how winner-loser effects determine later contest outcomes.

Early theoretical work regarding winner-loser effects primarily focused on mechanistic investigations of whether feedback from winner-loser effects explained the observed linearity of dominance hierarchies (reviewed in Lindquist and Chase 2009; Mesterton-Gibbons et al. 2016)). These models (termed Type I by Masterton-Gibbons et al, 2016) generally used a linear shift in contest ability, in which after a contest, an individual's intrinsic win potential was modified, whether by some fixed amount (Bonabeau et al. 1996), as a proportion of prior ability (Dugatkin 1997), or as a function of the difference between individual and opponent ability (Hemelrijk 2000; Hock and Huber 2006). This "ability-feedback" approach continues to be used in more recent models (e.g. Kura et al. 2016; Hickey and Davidsen 2019). All of these models treat winner-loser effects as a change in intrinsic contest ability, and cannot account for winner-loser effects where intrinsic ability is unchanged. While these models are useful in exploring the impacts of (intrinsic) winner-loser effects on the linearity of networks and provide an important foundation on which subsequent work expanded, we need a model which could isolate the effect of self-assessment per se.

Another class of models (termed Type-II by Masterton-Gibbons et al, 2016), focuses on how the strategic rules used in social contests could give rise to winner-loser effects. These models tend to be more evolutionarily minded and rooted in game theory, in terms of finding the optimal and/or stable strategies for social contests. There are a range of decision rules used, for example matching behavior to past probability of winning (Van Doorn et al. 2003; Fawcett and Johnstone 2010), or by using an actor critic model match expected reward with contest outcome (Leimar 2021). These models provide clear and even plausible mechanisms to explain

winner-loser effects and other features of dominance hierarchies, however these models either do not explicitly model self-assessment, or if they do, they leave it unchanged. As with the "ability-shift" models already mentioned, updating a decision rule is a different model than modifying an estimated self-assessment, even if the behavioral output may be similar. These decision rules essentially act as heuristics to capture an implied shift in self-assessed probability of winning. To our knowledge, only one model (Mesterton-Gibbons 1999), explicitly treats winner-loser effects as a change in self-assessment, but this is limited to three possible estimates: [Naïve, post-win, post-loss], and intrinsic ability only factors into contest cost, not contest outcome. Although these "strategy-based" models are very powerful, none of them are designed to explore whether winner-loser effects and dominance formation can be explained solely by agents refining their self-assessment of their intrinsic ability.

Given our goal to explore winner-loser effects as a process of modifying individual's estimates, Bayesian updating is the most natural framework to use. In brief, Bayesian updating provides a precise mathematical formula to modify an existing estimate of the state of the world based on new information (McNamara et al. 2006). It thus operates entirely on changes in individual assessment. Existing empirical research suggests that many animals use Bayesian processes to make decisions in other contexts (Luttbeg 1996; McNamara et al. 2006; Olsson 2006; Okasha 2013). Despite being ideally suited to calculate self-assessment of individual ability following a contest (as noted by Whitehouse 1997), to our knowledge, there is no model that treats winner-loser effects as a process of Bayesian learning. Admittedly, Bayesian modeling can be challenging to implement (as noted by McNamara and Leimar 2020), particularly as the

complexity of the information space increases, however understanding the specific predictions of changing self-assessments using Bayesian updating could provide important insight into the actual strategies used by different animal groups. Given the potential power of this hypothesis, it is important to identify the predictions of a Bayesian model and understand how we might distinguish Bayesian updating from alternative mechanisms.

Here, we present an agent-based model of winner-loser effects—and dominance formation generally—as a process of Bayesian updating. In brief, individual agents with self-assessed contest ability compete to win social contests, after which each agent uses Bayesian updating to modify their estimate of their own contest ability, which they then use to determine how much to invest in future contests. The purposes of this model were 1) to explore the features and limitations of Bayesian updating as it relates to winner-loser effects, thereby showing whether such a framework is plausible; and 2) to identify testable predictions that could empirically distinguish Bayesian updating from alternative mechanisms of modifying self-assessments. To first show whether Bayesian updating recreates known empirical features of dominance behavior, we identified specific features of winner-effects, i.e., the observed recency biases of winner/loser effects and the higher relative strength of loser effects as shown in the empirical literature (Bakker et al. 1989; Chase et al. 1994), as well the tendency of dominance hierarchies to be linear and stable. To achieve our second goal of identifying testable predictions, we focus on measurable attributes of fights—such as fight intensity or duration—and compare these measures across staged contests where the agents use either Bayesian updating or other updating mechanisms. For this comparison, we model linear shifts

in (self-assessed) contest ability, similar to (Dugatkin 1997; Bonabeau et al. 1999), and a fixed estimate, in which no change in self-assessment occurs. We show that Bayesian updating can reconcile disparate empirical observations of the attributes of winner/loser effects, while laying out clear, testable predictions that can be used in future experiments to identify the mechanisms by which dominance is established.

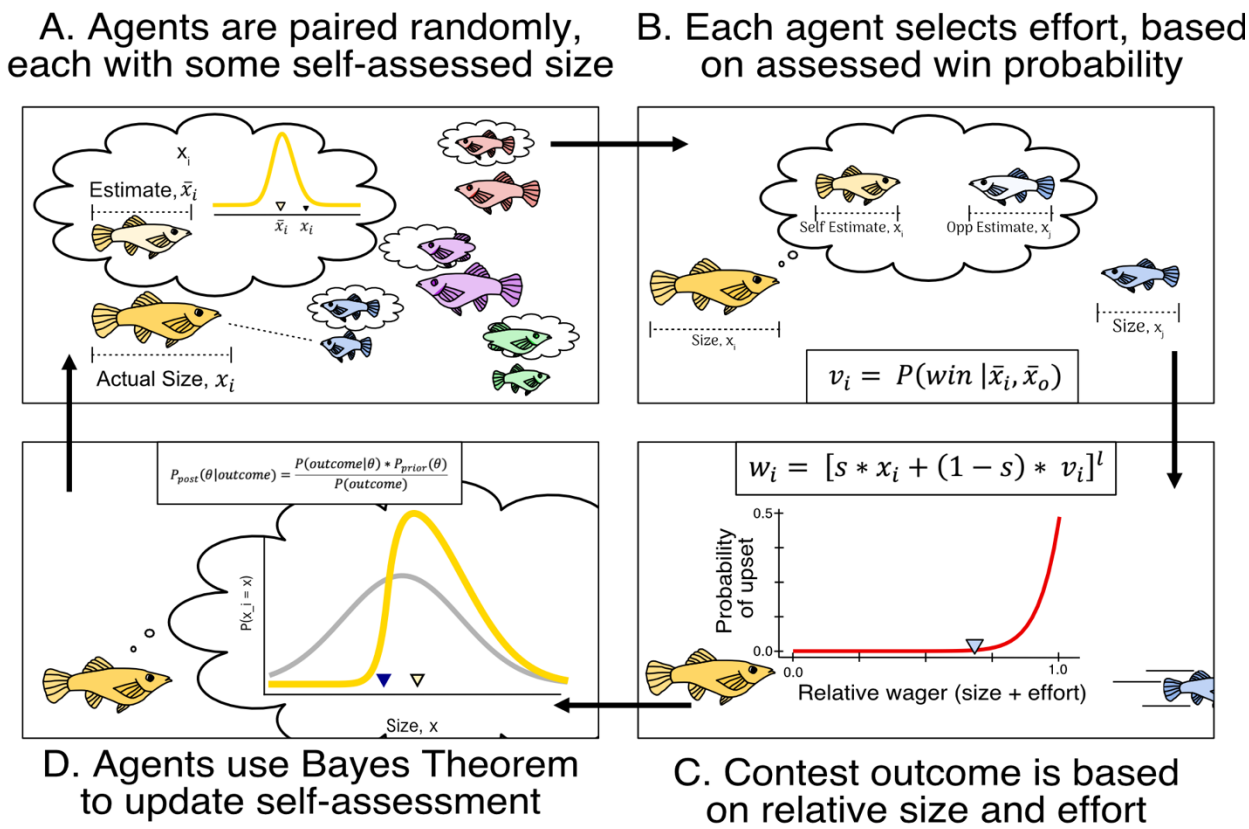

Figure 1: **A**, Individuals with fixed size (unknown to self) and an estimated self-assessment are paired. **B**, each individual estimates their own and opponent's size to determine their effort, based on their estimated probability of winning. **C (bottom right)**, the size and effort of each agent determines their relative wagers. The relative wager of the smaller individual determines the probability of winning (calculated as the probability of an upset by the underdog in a given contest). **D (bottom left)**, Based on

this probabilistic outcome, agents update their self-assessment, multiplying their prior assessment by the calculated likelihood of the observed outcome.

General approach

We envision a scenario where a group of animals compete for dominance in war of attrition contests (Maynard Smith 1974). As depicted in figure 1, we imagine these agents as fish of varying sizes, but individual agents could represent any species (mice, humans, etc.) competing in a social group. Our model consists of four steps, described graphically in figure 1. First, we generate a group of agents with random (normally distributed) intrinsic ability and imperfect self-assessments of that ability (figure 1A). For simplicity, we will refer to the intrinsic ability as "*size*", although in most natural systems other factors also contribute. Second, these agents are randomly paired and allowed to interact. In each contest, they select their effort, based on their estimated probability of winning, which is based on their existing self-assessment of their size and their assessment of opponent size (figure 1B). Third, contest outcomes are determined probabilistically based the relative size and effort of each participant (figure 1C, bottom-right). Finally, following the contest, each agent updates their estimate (i.e. their self-assessment) based on their prior estimate of size and their estimated likelihood of the observed outcome (figure 1D, bottom-left). This process is then repeated. Each step of the model is described in detail below.

Table 1 – Overview of parameters and their default values

Symbol	Default	Range	Explanation
x : $(x, x_i, x_o, \bar{x}_i)$	Random	$1 \leq x \leq 100$	Size of an individual. x_i represents actual size for individual i, while x_o is the actual size of opponent. $\bar{x}$ denotes an estimated size
μ_x	50	$1 \leq x \leq 100$	The mean of the population from which individuals are drawn
σ_x	10	$0 \leq \sigma_x < \infty$	Standard deviation of population from which individuals are drawn. At $\sigma = \infty$, uniform distribution is used.
v : (v, v_i, v_o)	None	$0 \leq v \leq 1$	Effort, i.e., the amount of time/energy invested in a contest. As with size, subscripts distinguish effort by the focal individual v_i from their opponent's, v_o
w : (w, w, w)	None	$0 \leq w \leq 1$	Wager, or contest performance, a combination of the agent's size and selected effort. This could also be thought of as work, i.e., the amount of energy imposed on opponent.
s	0.7	$0 \leq s < 1$	Relative importance of size vs. effort, acts as a simple multiplier when summing size and effort.
l	6.3*	$0 \leq l < \infty$	Relative wager dampener, essentially the importance of luck, acting as an exponent on relative wager, such that as l increases, the probability of an upsets decreases
σ_a	20	$0 \leq \sigma_a < \infty$	*self-assessment error*, the accuracy of individuals' starting guesses for their self-assessment.
σ_c	3.1*	$0 \leq \sigma_c < \infty$	*opponent-assessment error*, the accuracy of individuals when guessing the size of opponents prior to contests
$P(\theta)$	-	-	Used to denote the prior, an estimated, discrete distribution of the probability of being some size.

* For all unbounded parameters, we use a tangent scaling function for ease of use (see Supplemental), which is why these default parameters are arbitrarily specific.

Generation of agents

In order to model the development of self-assessment, we generate agents with some intrinsic win potential, called *size*, usually drawn from a truncated normal distribution between 1 and 100 (arbitrary units), such that

$$x_i \sim N(\mu_x, \sigma_x), \;\; where\ x_i \in [1,100] \tag{1}$$

Each agent has an estimate of their own size, $\bar{x}$, itself drawn from a truncated normal distribution centered around their actual size:

$$\bar{x}_i \sim N(x_i, \sigma_a), \;\; where\ x_i \in [1,100] \tag{2a}$$

where σ_a reflects individual *self-assessment error*, i.e., the initial precision with which they can estimate their own size. For all main results, $\sigma_a = 20$. Individuals then calculate a starting prior as a truncated normal distribution centered on their estimate, with standard deviation, $\sigma_\theta = \sigma_a$,

$$\mathrm{P_r}(\theta) = N(\bar{x}_i, \sigma_\theta), for\ 1 \leq x \leq 100 \tag{2b}$$

Effort strategy

Although the purpose of our model is to explore updating strategies, in order to do so, we must first establish how agents determine their *effort* in a contest, v_i, representing any degree of variable investment in a contest (e.g., the amount of time/energy an animal will expend before yielding.) Empirical observations (Hsu et al. 2008) and theory (Smith and Parker 1976; Enquist and Leimar 1983) suggest that an individual's effort should be proportional to their potential to

win, so we designed the effort strategy to match this. To determine their effort, each agent first assesses the size of their opponent, by drawing their estimate of opponent size, $\bar{x}_o$, from a truncated normal distribution centered on their opponent's true size,

$$\bar{x}_o \sim N(x_o, \sigma_c),\ where\ x_o \in [1,100] \tag{3a}$$

Here the standard deviation, σ_c, represents opponent assessment error, i.e., the precision of that assessment. For all main results, $\sigma_c = 3.1$ (see Appendix A for notes on parameter selection). Agents then set their effort, v_i, based on their estimated probability of winning, such that

$$v_i = P(win\ |\bar{x}_i, \bar{x}_o) \tag{3b}$$

where effort, v_i, is equal to the conditional probability of winning—given an agent's estimates about their own and their opponent's sizes and assuming both agents will equally invest 0.5 effort. This probability is calculated using the same function which determines contest outcome (see equation (4a), below.)

Determining contest outcome

For a given contest, the winner is decided probabilistically, based on the relative size and effort of each agent, combined to define each agent's *wager*, w_i:

$$w_i = s * x_i + (1 - s) * v_i \tag{4a}$$

Wager, w_i, and *effort,* v_i as terms are somewhat similar. It may be useful to think of wager in physical terms as *work*, i.e., the amount of energy imposed on the opponent, and effort as *vigor,* or the amount of energy expended by the individual. In other words, the wager scales the individual effort according to their size, where s allow us to control the relative importance of *size* and *effort*.

Because in our model (as in natural observations) contests are determined by when the subordinate individual yields, and contest success if a function of relative, not absolute, ability, it is convenient to calculate the outcome probability from the perspective of the lower-wagering individual, termed the underdog. We thus calculated the probability of an upset (a win by the underdog) with

$$P(upset) = \frac{1}{2} * \left(\frac{w_{min}}{w_{max}}\right)^{l} \tag{4b}$$

where l is the bias parameter, *luck*, controlling how close the two wagers need to be for there to be a meaningful probability of upset, such that as l increases to infinity, the probability of an upset goes to 0 unless the two wagers are equal. Once $P(upset)$ is established, we use a random number generator to simulate whether an upset occurs.

This approach for settling contests allows us to vary the relative contributions of size, effort, and stochasticity while observing the (simulated) behavior of individual agents and groups (e.g., how much individuals invests in contests, whether dominance relationships were stable, etc.). For all main results, we set the values of these bias parameters at s = 0.7, $l = 6.1$.

Updating self-assessment post contest

Here, we compare three possible strategies for updating self-assessment: Bayesian updating, linear updating, and static estimates (no updating), each described in detail below, in order of increasing complexity.

Static Estimate (No updating)

The simplest possible update strategy is to not alter your initial estimate. We model this scenario, where we keep each agent's initial estimate of size fixed throughout the simulation, regardless of contest outcomes. In the fixed case, an agent's estimate is their initial guess, x_i, with awareness, σ_a.

Linear Updating

A simple way to implement self-assessment updating, (modeled after Dugatkin 1997), is to either increase or decrease each agent's self-estimate of their size, using some proportion, p, of its prior estimate and/or adding or subtracting some constant, k.

$$\bar{x}_{i_{t+1}} = p\bar{x}_{i_t} + k \tag{5a}$$

To make this more comparable to our Bayesian updating approach, we set $k = 0$, and set limits on the maximum and minimum possible size ($X_{min} = 1, X_{max} = 100$), such that an agent's new estimate, $\bar{x}_{i_{t+1}}$, is a function of their previous estimate, shifted by some factor, where

$$X_{min} \leq \bar{x}_{i_{t+1}} \leq X_{max} \tag{5b}$$

We then allowed p to vary dynamically as a function of distance from the max/min size, in order to more closely replicate the variable shifts of Bayesian updating, where

$$p = \quad 0.1 * (X_{max} - \overline{x}_{i_t}), if\ win \tag{5c}$$

$$p = -0.1 * (\overline{x}_{i_t} - X_{min}), if\ loss$$

Since this function only acts on the point-estimate (defined by the maximum likelihood estimation of the prior distribution), after updating this max likelihood estimate, we generate a new, truncated normal distribution, centered around that estimate, based on their *self-assessment error*, σ_a, as in equation (2b). This approach provides a simple heuristic that, like Bayesian updating, allows for weighted updating, but updates self-assessment without calculating a likelihood function.

Bayesian updating

In contrast with simpler heuristics, Bayesian updating attempts to calculate the true probability of being a given size, based on known/estimated parameters about the state of the world and the probability of outcomes. In general, Bayesian updating modifies some *prior* estimate using a *likelihood function*. In our case, the prior is the agent's current estimate of its own size, modeled as a discrete probability distribution. The likelihood function is the probability of the observed outcome (i.e. winning a given contest), conditioned on being some assumed size, calculated across all possible sizes. Then the posterior distribution, i.e., the new probability distribution of an agent's size estimate, θ, is calculated by

$$P_{post}(\theta|outcome) = \frac{P(outcome \mid \theta) * P_{prior}(\theta)}{P(outcome)} \tag{6a}$$

We discuss the case of a winning outcome, but the formulation is similar following a loss. For each possible size in our discrete size range, the posterior probability of being a given size is

$$P(x_i = x|win) = \frac{P(win|x_i = x)*P(x_i=x)}{P(win)} \tag{6b}$$

where, $P(x_i = x)$ is the prior estimate of being a given size, while the likelihood $P(win|x_i = x)$ is the probability of an individual of size x winning the previous contest, which is calculated according to the probability of upset from equation (4b). $P(win)$ is calculated as the sum of $P(win|x_i = x) * P(x_i = x)$ across all sizes, x,

$$P(win) = \sum_{x=0}^{x_{max}} P(win|x_i = x) * P(x_i = x) \tag{6c}$$

After calculating the posterior estimate, agents update their point estimate of size, $\bar{x}_i$

$$\bar{x}_i = E(\theta) = \sum_{x=1}^{x_{max}} x * P(x_i = x) \tag{6d}$$

Where E is the maximum likelihood estimate of the size distribution. This value is then used to compute agent effort in the next contest.

Parameter space

Taken together, our possible parameter space comprises the contest outcome parameters (s, l), reflecting the influence of *size vs effort*, and *luck,* respectively; as well as the assessment parameters (σ_a, σ_c), representing *self-assessment error* and *opponent-assessment error*. For all main figures, we use the default values stated in Table 1, ($s = 0.7$; $l = 6.3$; $\sigma_a = 20$; $\sigma_c =$

3.1), which were chosen to approximate empirical observations (Beacham 1988; Bierbach et al. 2012), namely that upsets were rare, and contests were usually predictable based on size (figure A1). We explore the full parameter space is in Appendix B.

Model Analysis

The goal of our model is to answer two questions. 1. Does Bayesian updating for self-assessment produce dominance hierarchies that match the behavior of empirical systems? 2. How can we distinguish, empirically, Bayesian updating from alternative models of dominance establishment and winner/loser effects? To answer these questions, werun simulations of model behavioral experiments while extracting behavioral metrics that would be tractable in and relevant to an empirical context, such as contest intensity. These experiments, and the behavioral metrics we measure, are discussed in their relevant results sections.

Implementation

All code was written in python. The random number generator for contest outcome was from the *numpy* library. All code and results are available on github (github.com/aperkes/BayesBattleBots).

Results

Part 1: Matching Empirical Observations

1.1 Bayesian updating predicts the existence of winner and loser effects.

We first test whether agents using Bayesian updating show winner-loser effects. To do so, we model a simulated experiment (based on standard empirical approaches) by forcing either a win or loss in a contest between a focal agent and a “treatment” opponent of known size (in this case, the same size as the focal individual). We then measure winner/loser effects by pairing these focal winners (or losers) against “assay” opponents (a second size-matched agent). We repeat this process for 1000 focal winners and losers. We can then quantify the winner/loser effect as the proportion of focal agents which win against their assay opponent (where no winner-loser effect would yield 50% wins). As shown in figure 2a, we find obvious winner-loser effects, in that prior winners were more likely to win subsequent contests, while prior losers were more likely to lose. These winner effects are broadly observed across the range of parameter values (see figure B1).

<u>1.2 Bayesian updating does not drive dramatic differences between winner and loser effects</u>

Under these conditions, winner and loser effects are roughly equal in strength, with the loser effect being slightly stronger. Exploring the parameter space, although this loser-effect bias is broadly observed, these differences are very small (see figure B2). Although empirical results are mixed, this seems to conflict with known observations, in which loser effects tend to be dramatically more impactful. This discrepancy does not rule out Bayesian updating as a mechanism, since the strength of winner-loser effects depend strongly on the specific strategy used by an organism, which is in turn based on the costs and benefits of winning or losing. Given that we do not observe a strong loser-effect bias, it seems likely that the differences in the strength of winner and loser effects observed in natural systems come not from some

inherent bias in Bayesian updating per se, but from the strategies used by organisms to determine effort in response to changes in self-assessment.

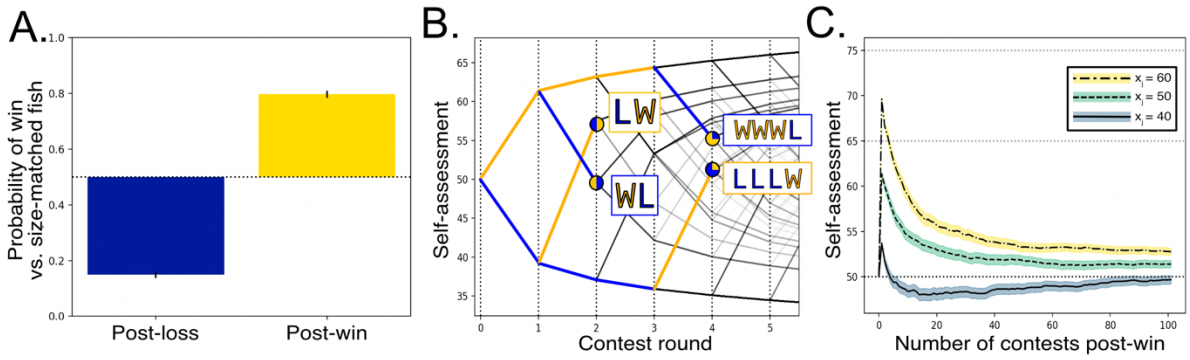

Figure 2: **A.** The proportion of focal agents (out of n=1000) winning against a naïve size-matched assay fish, after a treatment win/loss. **B.** Agent estimates are plotted over the course of repeated contests against naive, size-matched opponents, with the branching paths showing the estimates following a win or loss. The boxes note specific experiences highlighting the recency effect, which diminishes with increased experience. **C.** Focal agents' estimates change following a staged win against agents of varying size. The shift in the focal agent's estimate following a win is greater when their opponent is larger.

1.3 Bayesian updating produces a recency bias in winner-loser effects

Although winner-loser effects vary across species, there is strong empirical evidence that many species show a recency bias, in which the outcomes of more recent contests have a stronger effect on behavior than do the results earlier contests. We find that our model recreates these empirical observations. Once again emulating empirical tests of the winner-loser effect, we simulate contests between naive, size-matched agents and then pair these focal winners and losers against new, naïve, size-matched individuals, repeating this process over 6 rounds of

contests. We control the outcome of each contest to generate new branching paths with both winners and losers. To limit noise, all individual priors were centered on their true size, and agents' opponent-assessment was set to match opponent size exactly. Here, we measure winner-loser effects directly as the agents' own size estimates. In figure 2b, we see that winner-loser effects have an experience-dependent recency bias, such that "recent losers", (WL), who first won a contest and then lost, update their size estimate to be lower than "recent winners", (LW), who had lost and then won. This recency effect is based on the fact that the recent winners (LW) win while investing less in their second contest than the recent losers (i.e., initial winners, WL); the strength of this effect depends on the specific parameters used (see figure B3), but under all parameters, recent winners have posterior estimates that are greater than or equal to those of recent losers.

<u>1.4 Bayesian updating explains the observed variation in the persistence of winner/loser effects</u>

In empirical observations of dominance contests, winner-loser effects are persistent in some cases while in others they are short-lived or even non-existent. Here, we find that our model can reconcile these apparent discrepancies, based on the size of the initial opponent and the nature of agent experience between this "treatment" contest and the "assay" contest.

We again simulate a forced-win experimental approach, this time allowing a simulated focal agent of size $x_o = 50$ to win a staged contest against a treatment opponent of variable size ($x_o = [40,50, or\ 60]$). As in figure 2b, we measure the winner-effect directly, as the agent's own self-assessment. Following this initial treatment contest, focal agents are allowed to

interact with additional naïve individuals so that we could observe how their estimate changed over time. As we see in figure 2c, the scale of the initial shift depends on the size of the opponent, with larger opponents stimulating larger shifts in estimates. This initial winner effect attenuates rapidly with additional experience, although it does not fully return to baseline. How long winner-loser effects persist above a given threshold depends somewhat on the nature of these post-fight experiences (Supplemental figure B5), but even after 100 contests, the mean self-assessments of these initial winners are higher than their actual size, and those who won against bigger opponents have significantly higher self-assessments than those who won over smaller ones. Whether these changes could be detected is obviously a function of the power of the experiment used to assay winner-loser effects, however under Bayesian updating, the shift itself persists indefinitely.

Taken together, our model predicts that under Bayesian updating, the observed variation in winner effect duration (and whether it is detected generally) can be ascribed to differences in how the winner effect is generated, the experience of individuals in the interim, and how it is assayed in the subsequent test matches. All of this is consistent with empirical observations that the experimental methods used influence whether winner-loser effects are observed (Huang et al. 2011).

1.5 Bayesian updating produces efficient dominance hierarchies

Winner-loser effects are important not just for individual behavior but also the formation of dominance hierarchies. In natural systems, low intensity contests are a hallmark of natural

systems, as individuals seek to avoid serious injury in social contest. We find that our simulated agents using Bayesian updating recreates well documented features of natural dominance hierarchies (Jackson and Winnegrad 1988), by rapidly reducing the intensity of contests (figure 3A). This occurs in our model as the smaller agents begin to recognize when they are unlikely to win and yield more easily, resulting in lower mean effort across contests. This shift largely attributable to the rapid increase in accuracy shown in figure 3A, and is widely observed regardless of the specific parameters used (see supplemental figures B6 and B7).

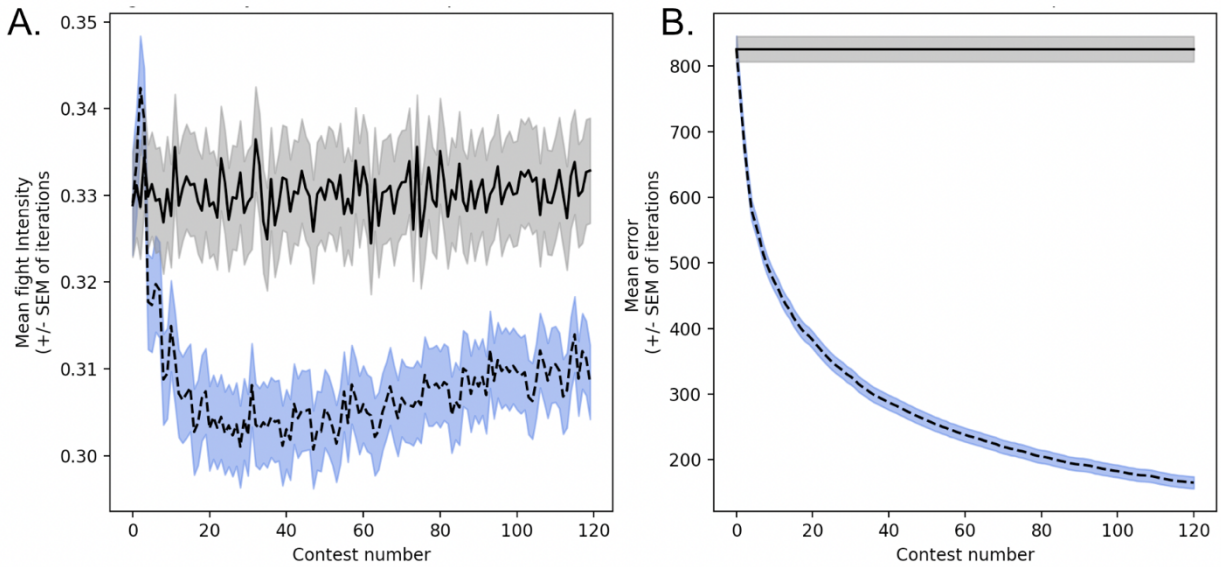

Figure 3: Emergence of stable dominance hierarchy through Bayesian updating. **A. Contest intensity decreases over time**. Here the mean intensity of fights, measured as the effort of the losing individual, decreases with repeated contests as individuals become more confident in their assessments. **B. Estimate error decreases over time.** Each line represents the mean stability across 1000 simulations for different updating strategies. Stability is measured by binning the simulation into 3-round windows, and comparing the consistency of dominance rank across bins. In both plots, blue dashed line shows Bayesian updating, while gray shaded solid line shows no-updating.

This improved accuracy also increases the stability and linearity of networks (see supplemental figures B8 and B9). Although network stability and linearity depend on the relative importance of size and other parameters, under most conditions, Bayesian updating produce networks that are more stable and more linear than those in which individual estimates are fixed, although it did not produce the hyper-stratification observed in models that directly modify intrinsic contest ability (e.g., Bonabeau et al. 1996; Dugatkin 1997).

Part 2. Distinguishing Bayesian updating from other models

We have shown that our Bayesian model is successful at recreating many empirical observations of animal contests, consistent with the hypothesis that Bayesian updating may be the mechanism behind winner-loser effects and hierarchy formation. To provide an approach that can empirically test whether animals are using Bayesian updating, as compared to other potential updating mechanisms, we explore two testable predictions of Bayesian updating, based on the two defining features of Bayesian learning—the likelihood-based updating and the informed prior. We investigate each of these two aspects separately, with simulated experiments to model future approaches in animal systems, focusing on measurable contest outcomes within empirically tractable scenarios.

<u>2.1 The effect of opponent size, but not focal agent experience, distinguishes Bayesian updating</u>

First, we focus on three potential updating mechanisms: Bayesian, linear updating, or static estimates (no updating). To test how these different updating mechanisms might be

experimentally distinguished, we simulate contests of focal individuals against different sizes and then measured the strength of winner-effects using a size-matched assay opponent match. As was shown in figure 2c, agents using Bayesian updating show stronger winner effects when their initial opponent is larger, while this “discrepancy effect” is not seen if the agents use a linear shift (figure 6a). This was also true for the loser effect (supplemental figure B10) and is observed wherever size and effort affect contest outcome (supplemental figure B11).

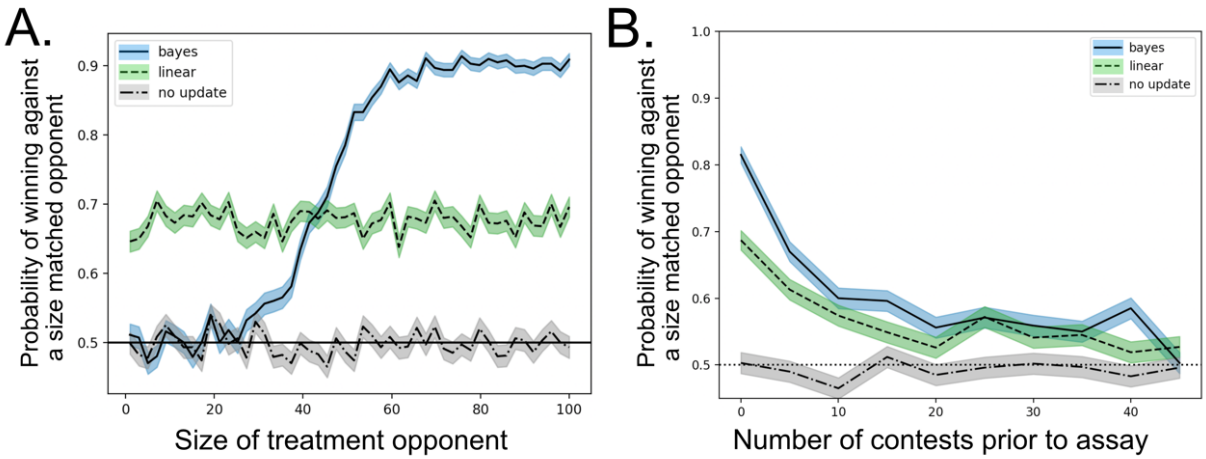

Figure 4: Simulated experiments show that updating strategy can be distinguished by the effect of opponent size and focal agent experience. **A.** Under Bayesian updating, the strength of the winner effect (the proportion of agents winning a size-matched contest after a staged win), is a function of opponent size. **B.** The strength of the winner effect is also a function of the number of contests prior to the forced win, however this does not distinguish Bayesian updating from linear updating.

It seems natural that variable winner-loser effects, based on the amount of prior experience, would also be a distinguishing feature of Bayesian updating, since Bayesian priors encode past

experiences and can reflect the confidence of an estimate, while linear updating has no such mechanism. To explore this relationship, we prime simulated focal agents with between 0 and 50 contests against naïve opponents (without fixing the outcome), thereby creating variation in the extent of each focal agent's experience. We then assess the strength of the winner-loser effect as before, first pairing the naïve or experienced focal individual with a size-matched 'treatment' opponent in which we force a win/loss, and then recording whether they win against a naïve size-matched assay opponent. As expected, in agents using Bayesian updating, winner effects are weaker for experienced agents than for naïve agents (figure 6b, dot-dashed line). Surprisingly, the same is also true for experienced agents using linear updating schemes (figure 6b, dashed line). This appears to be driven by the fact that under linear updating, estimates tend to diverge over time, becoming either extremely large or small, and are thus resilient to a single outcome. Thus, while this "experience effect" is a core prediction of Bayesian updating (across a range of parameter values, as shown in supplemental figure B13), and failing to find one would exclude Bayesian updating as a mechanism, finding an experience-attenuated winner effect does not rule out other possibilities. The effect of opponent size is the most straightforward single test for Bayesian updating.

2.2 Informed prior estimates drive consistent individual and context dependent differences

Often when behaviorists test for whether animals are using Bayesian strategies, they are asking specifically whether animals have "prior opinions" of the general state of the world which are then applied to each new context, for example the quality of a novel potential mate or some new foraging patch (reviewed in McNamara et al. 2006). In this sense, the question is not so

much about how individuals update their prior estimate during the task, but whether animals come into the task having informed prior estimates which reflect their expected state of the world, either based on their own past experiences or evolved over generations. With this in mind, we simulate experiments that would allow us to empirically distinguish whether individuals have informed priors, and how these priors form.

To explore this question, we simulate agents that begin with either informative or uninformative priors. As a reminder, priors are modeled as a probability distribution of size. For uninformative priors, we use either 1) a uniform distribution ("uninformed") or 2) a normal distribution centered on a uniformly random variable ("misinformed"); the informative priors are similarly divided into 3) a normal distribution centered on the agent's approximate size ("self-informed") or 4) a distribution that reflects the distribution of nearby agents ("peer-informed", for which we set peers to be either all large or all small)

We conduct simulated experiments in which we measure contest intensity—which in an empirical context could be quantified as the time spent in aggressive behaviors (Schlinger et al. 1987; Francis 1990)—among individuals using these different prior estimation methods. Each simulated agent is assayed 3 times, each time against a different, naïve size-matched opponent (drawn from the same prior method group). Because we simulate 3 trials for each agent, we can measure both persistent individual differences and context-dependent variation in the amount of time agents spent before yielding (i.e., *effort*). Here we prevent agents from updating their estimates after each contest. While allowing them to update estimates did not

change the qualitative results for the self-informed group, allowing other groups to improve their self-assessment during the contest resulted in the tendency to also correlate their effort with their size over time, although this correlation was quite weak (see supplemental figure B14).

In these simulated experiments, we find key differences in aggressive behavior which successfully characterizes groups based on the source of their prior (figure 7). Specifically, in groups of individuals using either self-informed or mis-informed priors, there are consistent individual differences in contest effort (i.e., significant variation within the group based on individual). Peer-informed agents may show individual differences if these agents were drawn from different peer groups (small peers vs large peers), but do not show differences when drawn from the same peer group (small peers vs small peers). Peer-informed agents (drawn from different groups) show within-group variation in effort correlated with peer size, but only self-informed agents displayed a correlation between individual variation in effort and size. This effect was observable across all parameters in which size was a factor in contest outcome (supplemental figure B15). Thus, by controlling for peer group and characterizing consistent individual variation in effort (and size), it is possible to identify whether individuals form prior estimates and, if so, what information they use to inform these estimates.

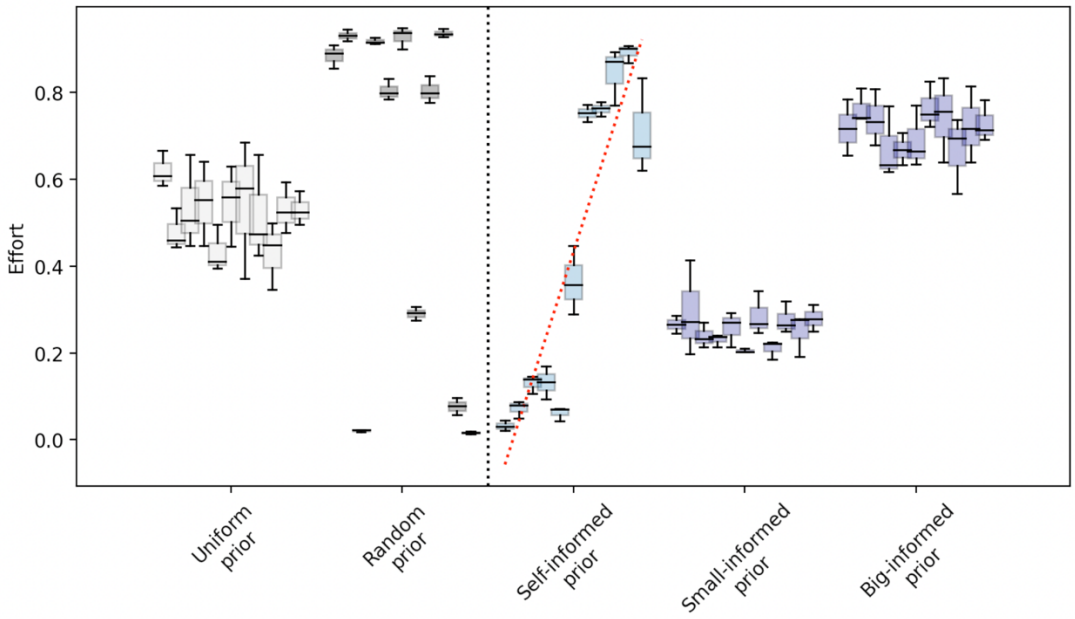

Figure 5: The y-axis shows contest intensity (effort) for 5 different simulated groups, each using a different starting prior. For each boxplot, boxes are arranged from left to right according to size. The left-most two groups (light and dark gray) have "uninformed" priors, either with all individuals having a uniform prior, or each individual having their own randomly selected prior. In the three groups on the right (light or dark blue), individuals have "informed" priors, using either internal information (Self-informed) or external information (small-informed/big-informed, the rightmost two clusters). Only self-informed priors lead to a correlation between size and effort (red dotted line).

Discussion

Here, we show that Bayesian updating can explain the existence of winner and loser effects resulting in contest behavior and dominance hierarchies that are consistent with those observed within social systems in nature. Notably, winner-loser effects occur without any explicit change in strategy (e.g., "invest more after a win"), simply as a side effect of agents making their self-assessments as accurate as possible after each contest. More importantly,

our model provides testable predictions to evaluate whether species' behavior is consistent with Bayesian updating for self-assessment. Our main findings here lay the groundwork for future work, both empirical and theoretical, in understanding the biological drivers of social dominance and the winner and loser effects.

Our Bayesian self-assessment model matches some, but not all, features of dominance interactions and winner-loser effects. Bayesian updating of self-assessment alone is sufficient to drive winner and loser effects, and because winner and loser effects are affected by the experience of the focal individual and the size and effort of the opponent, Bayesian updating can also explain variation in the relative strength and persistence of winner and loser effects, but it does not explain all aspects of network formation, which are likely mediated through mechanisms like contest strategy and feedback on intrinsic ability, which we do not model here. For example, agents using Bayesian updating increase the stability, linearity, and efficiency of dominance hierarchies, but in our model, Bayesian-based hierarchies do not stratify into hyper-linearity unless there are large differences in agent size from the start. This contrasts the work of intrinsic feedback models (e.g., Dugatkin 1997), in which linearity could emerge in the absence of initial differences. For both our model and prior intrinsic-feedback models, linearity and stability are a function of the intrinsic ability to win fights ('*size*'). Because in our model intrinsic ability does not change, agents only form more accurate assessments of intrinsic ability.

Similarly, Bayesian updating does not appear to generate strong biases towards loser (or winner) effects against size-matched opponents. The empirically documented loser-effect bias would thus need to originate from the contest strategy (i.e. how *effort* is calculated), or the contest function (i.e., how fights are settled), which we also do not explore here, however prior models based in game theory, most recently Leimar, 2021, show that stronger loser effects can be selected for when evolving contest strategies, suggesting that Bayesian updating, paired with evolved effort strategies, would replicate this behavior.

Given the success of strategic models (e.g., Leimar, 2021 and Leimar & Bshary, 2022) in replicating many of the empirical features of dominance hierarchies and incorporating a greater range of behavioral information (social eavesdropping, individual recognition, etc.), it is important to address whether the added complexity of Bayesian modeling provides any benefit. First, shifting an estimated self-assessment is a fundamentally different mechanism than updating contest strategy. It may be that given sufficient complexity, either could be replicate any given behavior, but we still expect the correlates of these approaches to differ. For example, a shift in self-assessment would likely be generalized across contexts (e.g., with novel predators or environmental challenges) whereas updating a strategy might be more limited in scope. More specifically, Bayesian updating is the most accurate possible approach for modifying an estimate, and thus provides a theoretical best case for any assessment-based approach, while also providing specific predictions which are different from existing strategy models. For example, while most strategy-based models encode some discrepancy effect (i.e., adapt more strongly to unexpected outcomes), they do not display any notion of estimate

598 “confidence”, whereas under Bayesian updating, experienced animals become more confident
599 in their self-assessment and thus more resilient to winner-loser effects (as shown in figure 2c).
600 Whether this matches the behavior of any given species is an empirical question, but doing the
601 work to generate such predictions will motivate the experiments necessary to identify
602 mechanisms.
603
604 In short, our model uses Bayesian modeling to explore the effects of self-assessment updating
605 in isolation, which produces specific, testable predictions, and provides insight into the possible
606 drivers of social behavior. Obviously this Bayesian self-assessment updating would need to be
607 integrated with additional mechanisms (i.e., more complex effort strategies, direct feedback on
608 intrinsic ability, social eavesdropping, etc.) in order to fully explain and predict social behavior.
609 Future work could synthesize these models to observe how different mechanisms interact, and
610 how the added complexity changes the predictions of our model.
611
612 It is important to note that while researchers often discuss “the” winner effect, there are of
613 course multiple winner-loser effects that occur via different mechanisms and operate across
614 different timescales. We have discussed winner-loser effects here from the standpoint of self-
615 assessment, contest strategy, and intrinsic ability, but even within any one of these effects, e.g.,
616 intrinsic ability, we can imagine varying, overlapping mechanisms, such as resource use and
617 acquisition, rapid changes in metabolism, and increased skill. Given that winner-loser effects
618 are both a natural consequence of, and solution to, the common problem of optimizing contest
619 effort, it should be no surprise that the specific behaviors and mechanisms used for a given

implementation vary based on the features and evolutionary history of the given species. Bayesian updating may be a fundamental aspect of at least one mechanism of winner-loser effects, and by developing and testing Bayesian models for social behavior, we can identify the mechanisms behind not just winner-loser effects specifically, but social interactions generally, in order to better understand and predict the formation and maintenance of social networks.

Supplemental PDF

This supplement is organized to match the structure of the main paper, and is divided into three appendices:

Appendix A, which provides greater detail and rational regarding the methods;

Appendix B, which has sub sections matching each results section, generally exploring the parameter space and, where relevant, the sensitivity of the results, while providing greater insight into the results; and

Appendix C, which compiles all equations and parameters used for easy reference.

Appendix A: Extended Methods

The effort strategy was chosen for simplicity and to be proportional to odds of winning

Theory and observations suggest that animals should, and do, invest in contests as a function of their probability of winning. Even within this limitation, there are unlimited possible updating strategies at various levels of complexity, and empirical work suggests that contest strategies vary both across (Arnott and Elwood 2008, 2009) and within species (Briffa et al. 2015). We assumed a very simple strategy—investing proportional to a rough estimate of the probability of winning. This approach does not fully take into account the certainty of self- or opponent-assessments, nor does it seek to guess how much an opponent will invest, and there is no optimization for costs and benefits or bet hedging. Furthermore, there is no variation in available energy, motivation, boldness, etc. All of these are likely important factors in real effort strategies, however exploring effort functions is beyond the scope of our primary goal: exploring the how changes in self-assessment would impact behavior. Future modeling should

explore how different effort functions impact winner/loser effects and social contests within the context of Bayesian updating.

Additionally, because the probability of winning is a function of the relative size of the smaller individual, in equation 3b, the probability of winning $P(win\ |\bar{x}_i, \bar{x}_o)$ varies depending on whether an agent estimates itself to be bigger or smaller than its opponent.

$$v_i = \left(\frac{s\bar{x}_i+(1-s)*0.5}{s\bar{x}_o+(1-s)*0.5}\right)^l \qquad if\ \bar{x}_i \leq \bar{x}_o \tag{A1}$$

$$v_i = \left[1 - \left(\frac{s\bar{x}_o+(1-s)*0.5}{s\bar{x}_i+(1-s)*0.5}\right)^l\right] \qquad if\ \bar{x}_i > \bar{x}_o$$

where, s and l are the simulation parameters used to account for the influence of *size* and *luck* respectively. Although the actual probability of winning depends on both relative size and individual effort, with this effort strategy, agents must approximate the probability of winning by assuming that both they and their opponent will invest the same amount in the contest.

Contests were determined in a way that allows us to explore the impact of size and effort

As with agents' effort strategy, there are unlimited possible ways that size and effort could combine to determine contest outcome. Here we aimed to find an approach that was backed in empirical observations and theory, while also being tractable and efficient in the context of our simulation.

First, we had to decide how size and effort combine to determine contest performance, what we call a wager. We wanted to be able to parameterize the relative impacts of size and effort, so the simplest way to do that was a weighted sum of these two inputs, as stated in the Equation 4a from the main text:

$$w_i = s * x_i + (1 - s) * v_i \tag{4a}$$

It seems likely that in natural systems there is also some interaction effect, as well some non-linearity in the output, so a more realistic function might be a something like

$$w_i = \beta_o x_i^a + \beta_1 v_i^b + \beta_2 (x_i * v_i)^c \tag{A2}$$

however, since the contest function is not the focus of our model, the increase in complexity did not seem warranted. That said, in applying this model to any given empirical system, researchers should take care to ensure whatever function used is relevant to their system.

Next, we had to decide how agents' wagers should determine outcome. Empirical observations suggest an S-shaped curve for contest outcome as a function of relative size, that is, contests are highly predictable when individual size is dramatically different, and unpredictable when agent size is very similar, as shown in figure A1 below, reprinted from Graham & Angiletta, 2020.

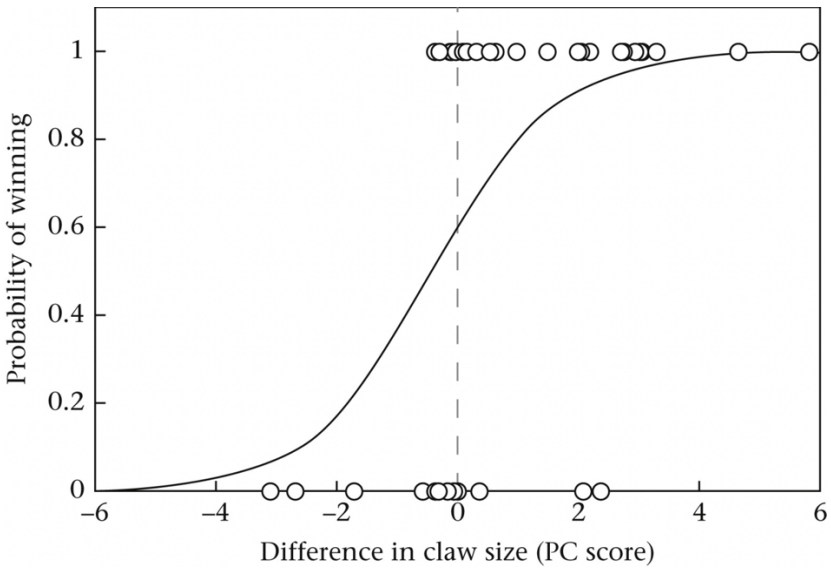

Figure A1: Empirical measure of the role of intrinsic attributes (claw size) in determining contest outcome, (from Graham and Angilletta 2020). In this case, overall body size was matched between participants, in order to observe the effect of claw size specifically. Note that while the sigmoid function used above is fit to the absolute difference in claw size, our model is based on relative size of the smaller individual.

As a theoretical explanation of this behavior, we hypothesize that there is some noise around an individual's expected performance—based on their size and effort—and the actual output. This noise would come from both intrinsic and extrinsic factors, such that each individual's actual wager, w_i, is drawn from a normal distribution centered on the ideal wager, $\overline{w}_i$, (based on their intrinsic potential, *size*, and their behavior, *effort*) (figure A2a). A contest represents drawing a single random value from each performance distribution, where the higher observed performance wins. If we implement this model to generate performance distributions for each individual from which we could draw their "output" wagers, we can then calculate the

likelihood of the lower performing underdog being the larger draw. Since these are two normal distributions, $P(w_i > w_j) = P(w_i - w_j > 0)$.

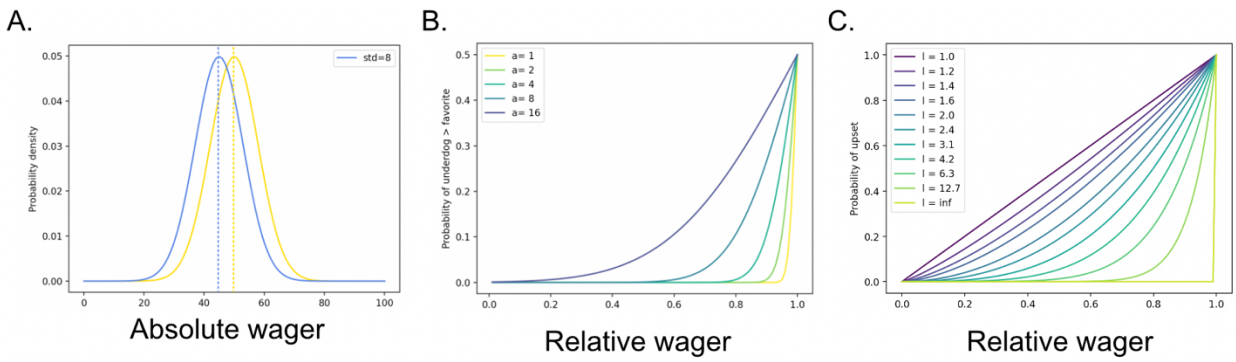

Figure A2: Contests are determined as a function of the relative performance of competing agents. A. Given some theoretical "output" performance by two different performing individuals (dotted lines), we assume there is a distribution of "observed" outputs. B. Given two probability distributions, each line shows the probability of upset as a function of the relative wager (of lower individual), which is the probability of drawing a larger value from that lower distribution. The shape of this function is determined by the standard deviation of the observed performance distribution. C. We approximate this process by raising the relative wager to some exponent, l.

This model has the parameter σ_w, the amount of variation in performance output. For different values of σ_w, we get likelihood functions that are more or less convex (figure A2b), representing contests that are more or less biased towards the higher-performing individual. Note that in contrast with the empirical figure A1, here we are quantifying the actual contest performance,

722 the combination of size and effort (which are abstractions of body size, weapon size, contest
723 skill, energy investment, time investment, etc.), rather than just size.
724
725 While this approach is justifiable both theoretically and empirically, it is computationally very
726 expensive to calculate the cumulative distribution multiple times for each contest, especially
727 when running millions of iterations of simulations. For this reason, we chose to implement the
728 simpler exponential described in equation (4a). This produces the win probability function and
729 parameterization shown in figure A2c, both of which are similar to the CDF-function, while
730 being both simpler and roughly 100x faster to run. In the main results, we set $l = 6.3$ (see note
731 on parameter values below) but explore the range of l values shown in figure A2c.
732
733 As a note, under this relative wager calculation the agent with greater relative size and effort
734 (i.e., the '*favorite')* usually won, but the agent with the smaller wager (the '*underdog*') could
735 win, particularly if the two wagers were close. It is important to highlight that the smaller wager
736 was not necessarily the smaller sized individual; if a relatively smaller agent contributed far
737 higher effort, they could have the larger wager, w_{max} , and thus be the '*favorite*' for the
738 purposes of our calculations.
739
740 Like our chosen effort strategy, potential contest functions are not the focus of the model, but
741 these model decisions have an influence on self-assessment updating and the winner effect,
742 since the likelihood function used when updating the prior estimate is calculated based on our
743 chosen contest function. Future work could explore how different contest functions impact the

behavior and predictions of Bayesian updating, and compare these to observed examples of social dominance.

Further details of Bayesian updating

During Bayesian updating, focal agents multiply their prior self-assessment (a probability distribution of their size) by the likelihood function of the observed output (winning or losing against the specific opponent). Finding this likelihood function requires calculating the probability of winning (or losing), conditioned on the focal agent being each possible size, x , across the range of possible sizes. Because probability of winning is calculated as function of the smaller relative wager, $\frac{w_{min}}{w_{max}}$, the probability of winning is equal to the probability of upset when the focal agents hypothetical wager is less than the opponent's wager ($w_i | \, x \leq w_o$), but is its inverse (i.e. probability of no upset) once the focal agent's hypothetical wager is greater than the opponent's:

$$P(win|x_i = x) = \; P(upset), if \; w_i | \, x \leq w_o \qquad \text{(A3)}$$

$$P(win|x_i = x) = 1 - P(upset), if \; w_i | \, x > w_o$$

Note that this inflection point is based not on relative size but on relative wager, such that agents must take into account both their own and their opponent's effort into their calculations of relative wager, and thus the probability of winning, $P(win)$. For simplicity, we assume that after the contest, agents have perfect knowledge of the size and effort of the opponent, however this calculation could be extended to allow for estimated distributions of opponent

size and opponent effort, assuming we have probability distributions for opponent size and effort.

Overview of parameter space

There are 4 core parameters (s, l, σ_a, σ_c) which define the main parameter space. For any given simulation, there are additional parameters, such as the population makeup or the experimental design (e.g., the size and effort of a staged opponent), which are discussed in their relevant sections. For parameters within a fixed range (e.g., s has the range [0,1]) we sample evenly across this range.

For unbounded parameters (e.g., $\sigma_a \in [0, \infty)$), we use scaling functions for ease of use in implementation. For l, we use the scaling function

$$l = \text{cotangent}\left(\pi * \frac{(1-L)}{2}\right), \quad 0 \leq L \leq 1 \tag{A4a}$$

Our default $L = 0.8$, corresponds to $l = 6.3138$. For simplicity, we show the scaled, l values in our plots, rounded to the first decimal place, which represent sampling evenly across the range of L values. This function can also be used to map L to values of l between $[0,1]$, however in our parameter exploration, these results all behaved very similarly to $l = 0$, so we focused on values between $[0,1]$, which are also far more relevant to our behavioral model.

Similar to l, for both σ_a and σ_c, we use the scaling function

$$\sigma_a = 20 \tan\left(\pi * \frac{A}{2}\right), 0 \leq A \leq 1 \quad \text{(A4b)}$$

For implementation purposes, evaluating at the limits sometimes requires some conditional statement to avoid errors, for example where $\sigma_a = \infty$, we use a conditional statement to set the assessment distribution to be uniform across the range of sizes.

Our default assessment parameters assume a situation where opponent assessment is always fairly accurate ($\sigma_c = 3.1$), while naïve self-assessment is very noisy ($\sigma_a = 20$). This is based on the idea that observing opponent size is actually easier than observing your own size, but this is not necessarily the case. We explore the full range of this space in the figures below, including where $\sigma = \infty$, in which opponent assessment is a random uniform variable, and $\sigma = 0$, where assessment is perfect. The case where $\sigma_c = \infty$ also functions as a version of pure self-assessment, albeit with a strong random component.

Appendix B: Exploration of Results

1.1 Bayesian updating drives winner and loser effects

First, it is useful to highlight the components which drive winner-loser effects in this model. Winner-loser effects can be broken down into 3 factors: the change in estimate, the extent to which that estimate leads to a change in effort, and the extent to which the change in effort leads to a change in outcome. Parameters can act on one or more of these mechanisms, but even changing only one mechanism can impact winner/loser effects (so, for example, changing how effort impacts outcome could strengthen winner effects, even if the change in effort and estimate remained the same.

Figure B1 repeats the simulation show in figure 2a across the range of possible parameters. In short, focal agents are selected to either win or lose against a size matched 'treatment' agent, and then tested against a second, size-matched 'assay' agent in order to measure the strength of the winner effect. For simplicity, starting priors always begin centered on the true agent size, $x = 50$. We then compare the change in effort estimate (pre-treatment vs post-treatment), the change in effort (treatment vs assay), and the assay contest win rate. We show that across all parameters where size contributes to contest outcome, winning increases estimated size, while losing decreases it. The only exception is where $s = 0$, in which wagers are entirely determined by effort, $l = 0$, where contests are entirely determined by luck, and $\sigma_a = 0$, in which agent's being with 100% confidence of their correct size.

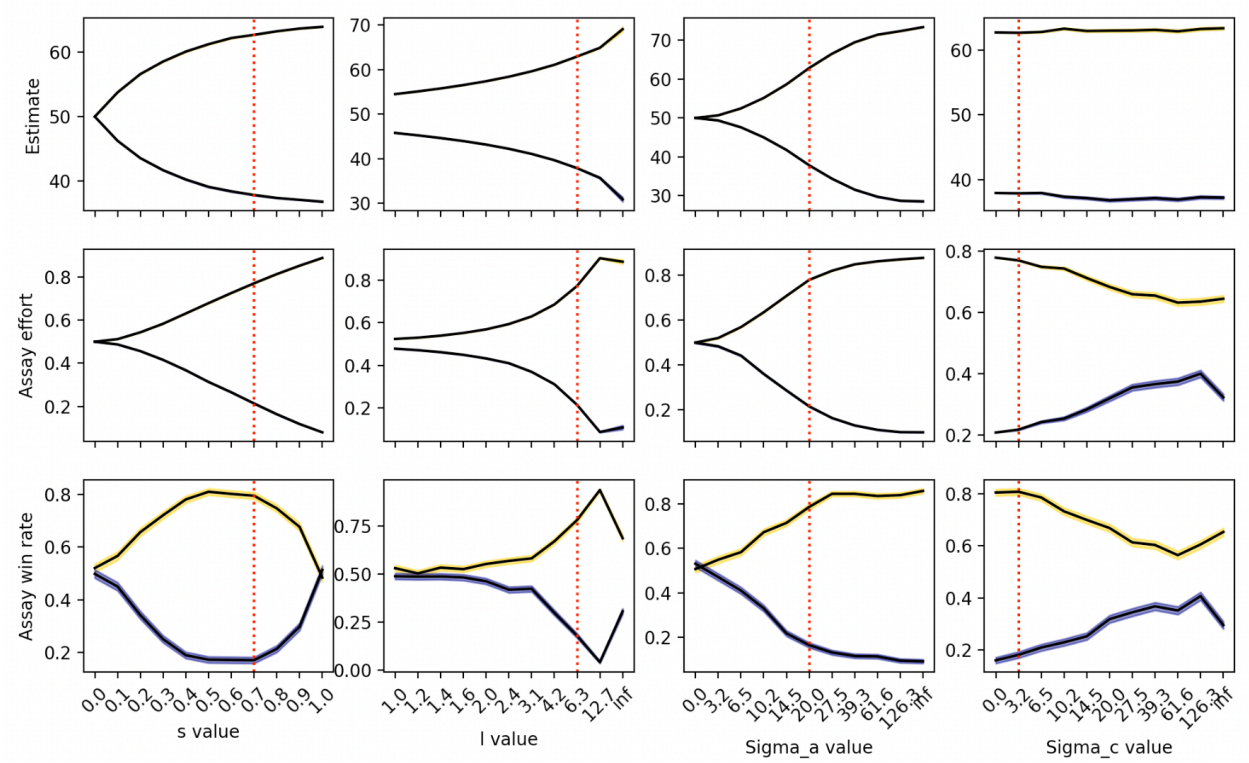

Figure B1: The strength of winner-loser effects across a range of parameters. How do estimates, efforts, and probability of winning change for different parameters. Sensitivity analysis across the whole parameter space for size matched contests. Lines represent the mean outcome from 1000 iterations, with shading shows the mean +/- 1 SEM. Gold lines show post-win, blue lines show post-loss. The variation comes from error in initial self-assessed estimates, and the error in opponent-assessment leading to variation in effort during both the treatment and assay contests.

There were additional parameter ranges where, despite a difference in agent estimate, there was no observable difference in win rate, notably wherever l was close to 0, since this flattens the probability of winning, resulting in similar outcomes. This also, naturally, occurred where $s = 1$, since a change in effort cannot impact the outcome where wagers are entirely determined by size.

The shift in estimate is also a function of the opponent used, as shown in figure 2a. In the above figures, we used an opponent of equal size and effort to the focal fish, however winning against a large fish with high effort, will of course have a more dramatic impact on estimates than winning against a small or low-effort fish. Similarly, losing to a small or low-effort fish is very impactful, while losing to a large fish has little change on estimate.

1.2 Relative strength of winner and loser effects

As with the *presence* of winner/loser effects, their relative strength, compared to one another, is mediated by changes in estimate as well as how that impacts effort and, in turn, contest outcomes.

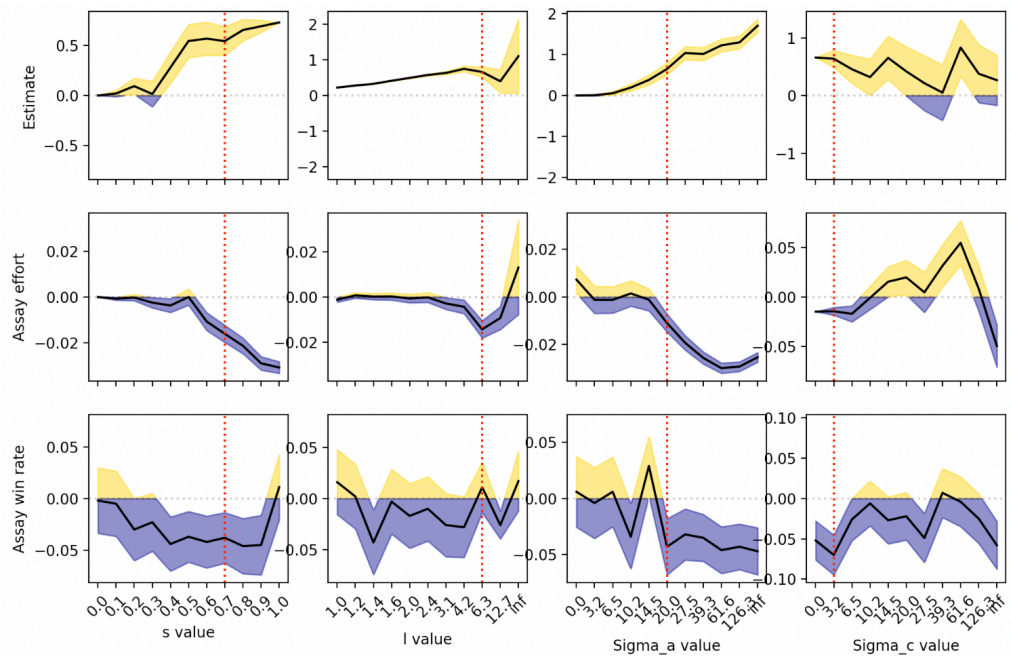

Figure B2: There is little or no bias towards winner or loser effects under Bayesian updating. This simulation for this plot is identical to figure B1, however here we plot the difference in the strength of the winner and loser effects (i.e., the absolute difference from the null value, either 50 or 0.5). The color here simply shows whether the outcome is biased towards winner effects (gold) or loser effects (blue).

Over most parameters there is a slight bias towards increasing estimates (i.e., a stronger winner effect). This appears to be due to how the likelihood function is based on relative size, and thus flattens out with increasingly large estimates. In comparison, the effort function amplifies loser effects somewhat, particularly where size is important and initial self-assessment error is high. This also seems to be based on the importance of relative size, since increasing an estimate by some fixed amount (e.g., 25 units) results in a smaller perceived relative change (50 : 75, where $\frac{w_{min}}{w_{max}} = 0.66$) than decreasing your estimate by the same amount (25 : 50, $\frac{w_{min}}{w_{max}} = 0.5$). Taken together, in most cases, winner and loser effects have a similar impact on win rate, with slight bias towards the loser effects over most parameters.

It is also important to note, even in the cases where there is a significant difference in winner vs loser effects, this bias is very small (the maximum difference observed is only around 5%). In empirical examples, winner and loser effects can be dramatically different, and in some examples only winner (or loser) effects are observed. Within our model assumptions, these outcomes are unlikely to occur, without altering the size or effort of the agents (for example, having the 'treatment' opponent be very small would cause a strong loser effect but no winner effect'), however empirical tests generally use size-matched contests, so this is unlikely to explain the observed strength of the loser effect. It seems likely then that the majority of differences in winner and loser effects are driven either by the specific effort strategy used, or else in how contests are determined. As an illustrative example, the effort strategy "invest 100% whenever there's >40% chance of victory, otherwise invest nothing" would create a

strong loser effect and no winner effect. These strategies are likely driven by the specific costs and benefits of competition, which are different for each taxa, and those seeking to make predictions for specific systems should seek to identify the relevant underlying functions for that specific context.

1.3 The recency effect

Initially, the recency effect shown in figure 2b may seem to conflict with Bayesian updating, which is known to be order invariant, however this occurs because earlier fights determine self-assessment, and thus effort invested during later fights. Winning first means that an agent invests more in the second fight. Because prior-winners invest more, the likelihood of losing is lower except where size is very small. This results in prior winners being more susceptible to loser effects, while prior losers are more sensitive to winner effects.

To assess the sensitivity of the recency effect, we use a simplified version of the simulation in figure 2b. Again using size-matched agents, we compare two conflicting cases, recent winners (LW) and recent losers (WL), and compare the final estimate, effort, and outcome in an assay contest.

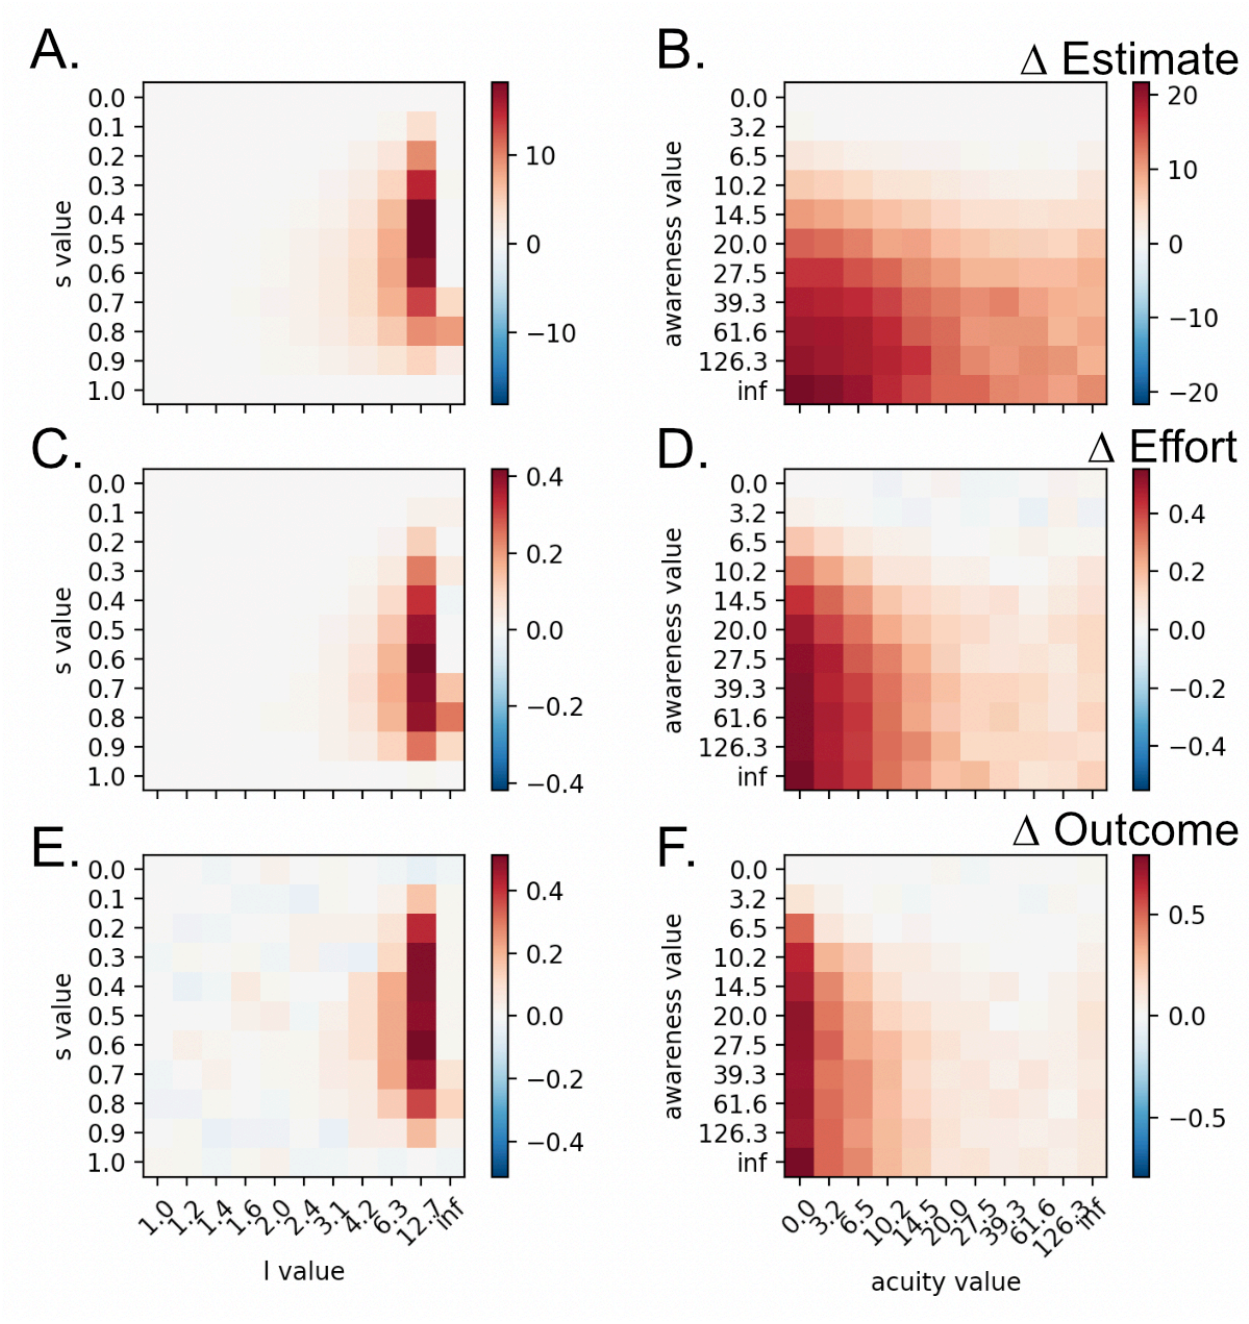

Figure B3: Here agents experienced two contests with conflicting information, i.e. a win then loss (WL) or a loss then win (LW). Heatmaps show the difference between these outcomes. The top row, A-B shows the difference in final estimate, showing recent winners were generally higher (red). The second row (C-D) shows that the same is true for invested effort. The bottom row, E-F, shows the results of a size-matched assay

contest. Under all parameters, the recent winners are (statistically) more successful than or equal to recent losers.

There is a consistent bias towards more recent outcomes, in this paradigm where the recent event conflicts with the prior event. This is particularly true where both size and effort play a role, and where contests are strongly biased towards the higher-wagering 'favorite'. Where s is small, or large, the difference between recent winners (LW) and recent losers (WL) disappears. Similarly, where self-assessment error is very low, or opponent-assessment error is very high, this effect disappears. There are no scenarios where we observe a significant primacy effect, i.e. the first event overrules later conflicting information.

That said, this recency effect is not necessarily strong enough to overcome any history: with each additional experience, the difference decreases until eventually the one e.g., win is not enough to supersede the sum of losses. Across all parameters tested, it never took more than 4 prior events (e.g., 4 wins) to overcome the recency effect, such that WWWWL always results in higher win-rate than LLLLW.

1.4 The strength and persistence of winner/loser effects

This work was motivated in part by seeking to reconcile conflicting empirical results as to the strength and persistence of winner-loser effects. For empirical work, the persistence of the winner effect is really a question of whether experimenters are able to *observe* a winner effect using some empirical test. There are many potential ways to test for winner effects, but for

simplify, we will only explore this question in the standard experimental assay ('treatment' win/loss followed by an 'assay' contest).

As we show in Figure 2c, this strength and persistence of winner effects is a function of the initial event, the intervening experience Figure S8 expands this result for both winner and loser effects.

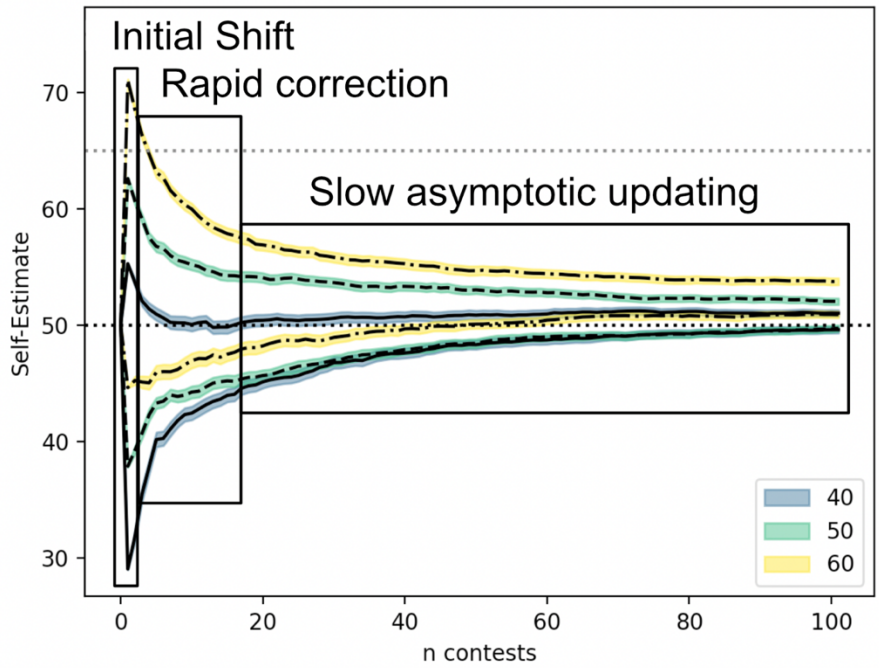

Figure B4: The persistence of winner and loser effects are a opponent size and intervening experience. Here focal agents were subject to a single treatment contest (either win or loss) followed by repeated, unbiased contests against other agents. The long-term behavior of focal agents following a forced win/loss is characterized by an initially large shift, followed by rapid correction (which can sometimes overshoot the correct value, as in the blue-winner line), and then a long slow process that asymptotically approaches the true value.

The initial change in estimate is simply a question of winner/loser effect strength (figure B1). The greater this initial change, the longer it will last. This is a natural prediction of Bayesian updating, and contrasts with 'one-shot' models. In addition to this initial shift, the experiences agents have between the 'treatment' contest and the 'assay' are also critical in determining how long winner-loser effects persist (figure B5).

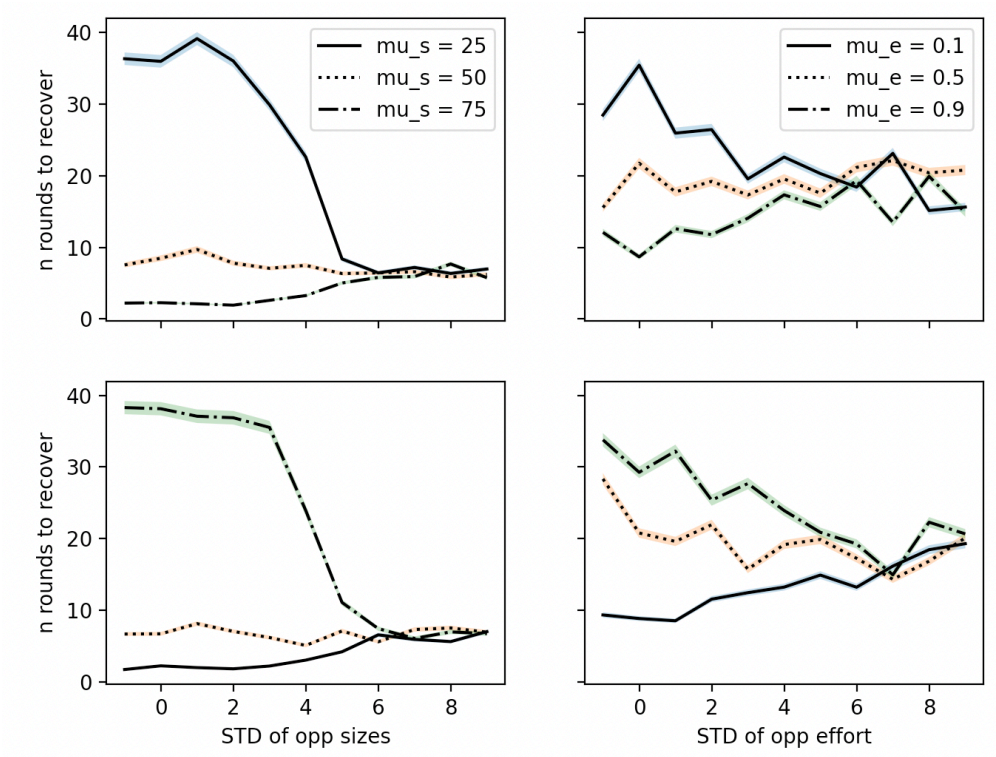

Figure B5: The top row (**A-B**) shows the time to recover from a win. Recovery is the number of rounds needed to return to within 10 units of actual size (see figure B4). Here each round represents a new individual with randomly selected size and effort. **A.** shows the number of rounds to recover with respect to the standard deviation of opponent sizes, plotted for three different mean sizes, while **B** shows this as a function of the standard deviation of opponent effort, again plotted over 3 different mean efforts. **C** and **D** show the same metrics after a staged loss.

We show in figure 2c that winner effects attenuate over time, but the speed of this attenuation (i.e., the number of rounds required to recover) depends on the diversity of info, or more precisely, the frequency with which an agent receives information which conflicts with its current self-assessment, as shown in figure B5. Thus with a greater variety of opponent sizes and efforts, agents can more quickly return to near-accurate self-assessments. Similarly, with larger mean-sized opponents will quickly erase a winner effect, but prolong a loser-effect (see green dash-dotted line in figure B5a)

1.5 Accuracy, Efficiency, Linearity, Stability

As shown in figure B4, under Bayesian updating winner/loser effects are essentially temporary misinformation, in that either estimates were wrong and are now better, or that they were better and are now worse. Once agent estimates are perfectly accurate, agents will perform as expected by their size. Thus, in our model, as agents approach accurate assessments, network metrics like linearity and stability become a function of intrinsic differences and contest parameters. Given this, it is perhaps not surprising that under most conformations, we do not observe hyper-stratification. In other words, under Bayesian updating, perfectly stable, linear hierarchies do not usually emerge from groups of similar or identical individuals.

We do however observe that Bayesian updating generally has a stabilizing influence on networks, compared to fixed, naive estimates. As in figure 3, we simulate groups of agents (n=5) and allow them to interact over 10 rounds (40 contests per individual), simulated over 100 iterations, from which we calculate group statistics (discussed below).

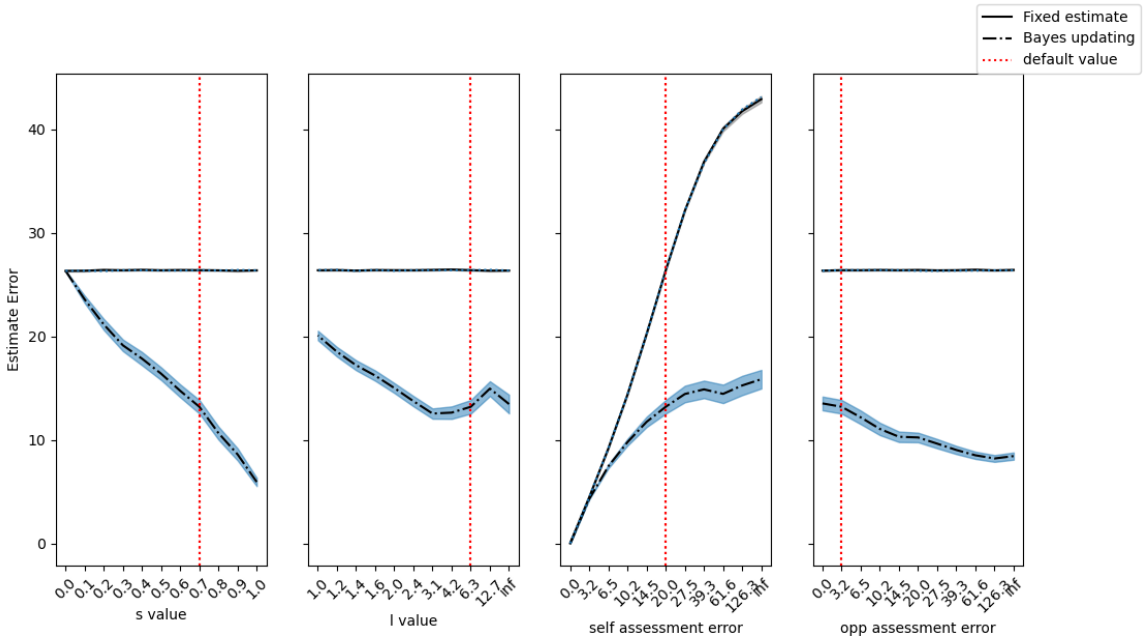

Figure B6: Estimate error plotted over a range of parameters. Groups of 5 agents were allowed to interact for 10 rounds, and lines represent the mean error (+/- SEM) over 100 iterations. The blue line represents the final estimate error for Bayesian updating agents, while the gray line shows the final error for groups with no updating. The dashed line shows the first round estimate error for Bayesian updaters. Estimate error is calculated based on squared difference between actual size and estimated size +/- prior standard deviation, as shown in equation (B1).

To calculate accuracy, we measure the distance from the true agent size, x_i and the $\bar{x}_i \pm \sigma$ confidence interval,

$$\varepsilon = \sqrt{(\bar{x}_i + \sigma_\theta)^2 + (\bar{x}_i - \sigma_\theta)^2} \tag{B1}$$

Under almost all parameters estimate error decreases compared to initial estimates (the exception being the edge cases where size has no impact on contest outcome, i.e., $s = 0$ or $l = \infty$). This does not occur in agents that do not perform updating (figure B7, gray shaded line)

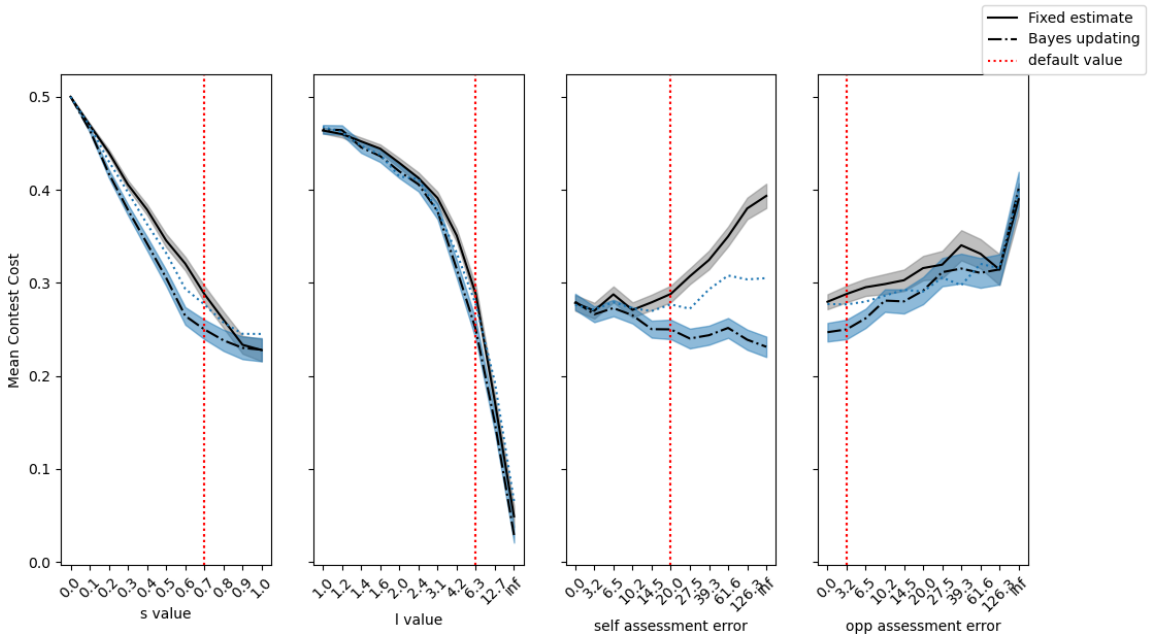

Figure B7: Contest cost plotted over a range of parameters. As in figure B6, groups of 5 agents competed over 10 rounds. Lines represent the mean cost (+/- SEM) over 100 iterations. The blue line represents the final-round mean cost for Bayesian updating agents, while the gray line shows the cost for groups with no updating. The dashed line shows the mean first round cost for Bayesian updaters. Contest cost is equal to the effort of the lower-investing individual.

Similarly, contest cost, measured as the effort of the lower-investing agent, decreases over the course of the simulation, however the effect of the Bayesian updating is much smaller than the effect of the simulation parameters themselves, i.e., contest cost is low when l is high, such that contests are highly biased against the 'underdogs', since an individual will invest proportional to their probability of winning. Where size is unimportant, or where initial

assessment is high, we do not observe the emergence of low cost behavior that is typical of many empirical dominance hierarchies, suggesting that Bayesian updating provides the greatest benefits where naive self assessment in poor and where size and effort both contribute to contest outcomes.

Like reduced cost contests, linearity (figure B8) and stability (figure B9) are hallmarks of social dominance hierarchies, and we find that under most conditions where contests are biased against the underdog (i.e., $l > 6$), Bayesian updating leads to a slight increase in both linearity and stability.

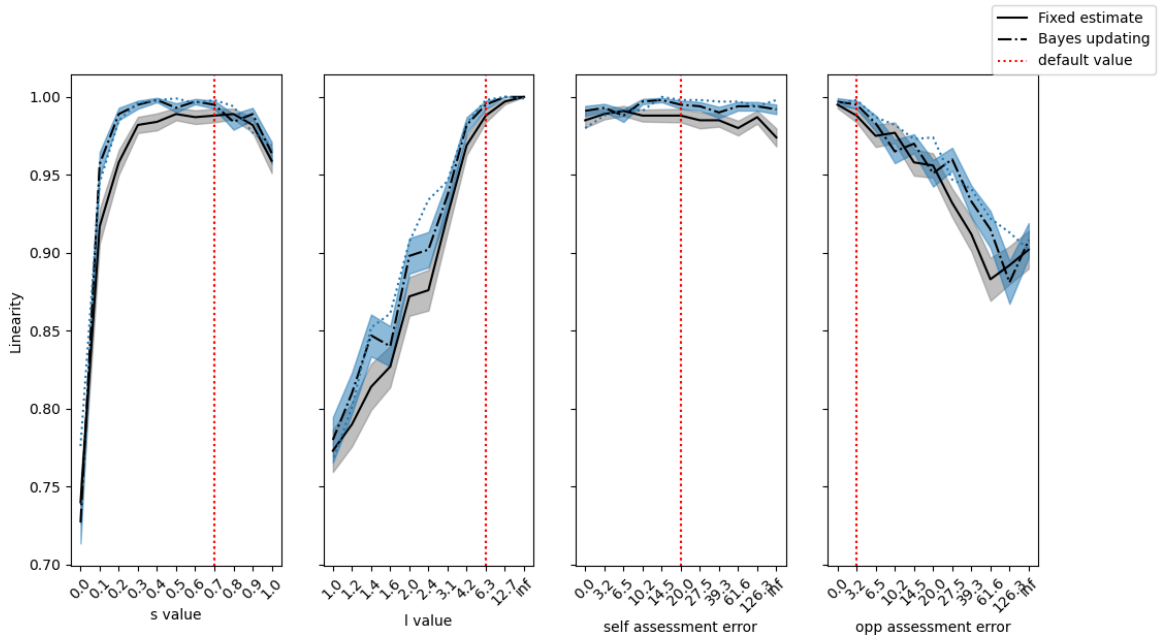

Figure B8: Network linearity plotted over a range of parameters. As in Figure B6, groups of 5 agents competed over 10 rounds. Lines represent the mean, final-round linearity (+/- SEM) over 100 iterations. The blue line represents the final-round linearity for Bayesian updating agents, while the gray line shows the

linearity for groups with no updating. The dashed line shows the mean first round cost for Bayesian updaters. Linearity is measured as the proportion of triads which are transitive, following Appleby's approach.

To measure the effect of Bayesian updating on the linearity of dominance hierarchies, we calculate linearity for the first and last rounds using Appleby's approach (Appleby 1983), measuring the proportion of triads which are transitive. (if A dominates B, and B dominates C, A dominates C). Stability is calculated by composing per-round dominance hierarchies (a round represents a series of contests in which every agent faces every other individual). We then bin the first 3 rounds, and the last 3 rounds, in which we measure the proportion of pairwise dominance relationships which are consistent for all rounds within a given bin (for two agents, if A dominates B in round 1, round2, and round 3, that is consistent, otherwise, it is inconsistent).

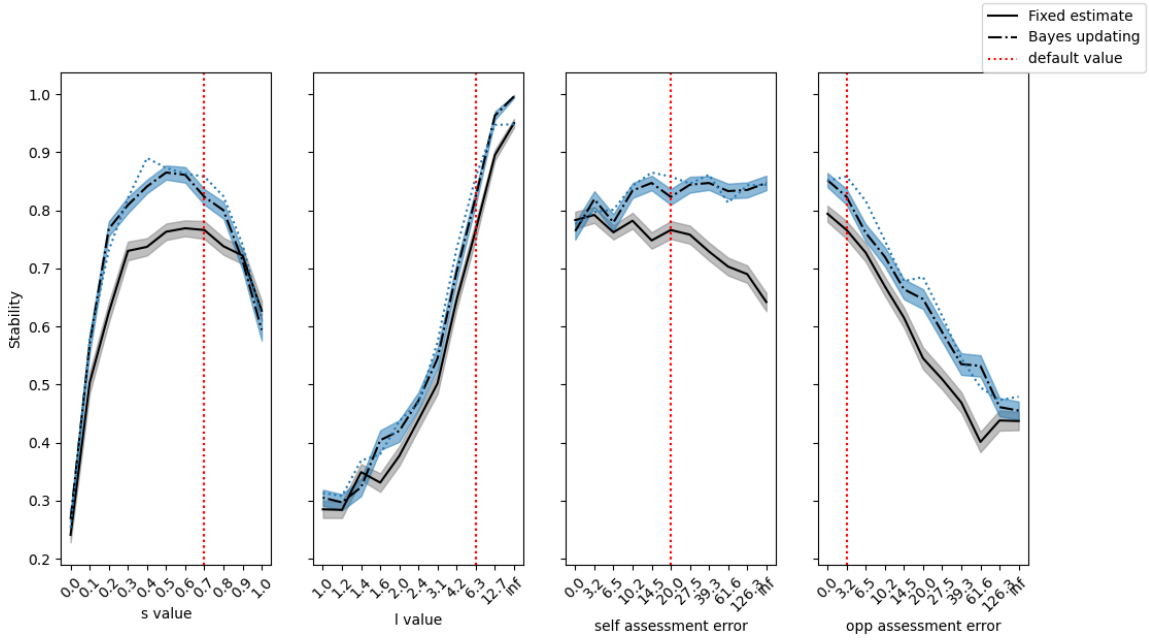

Figure B9: Network stability plotted over a range of parameters. As in Figure B6, groups of 5 agents competed over 10 rounds. Lines represent the mean stability (+/- SEM), calculated over the final three rounds of the

contest, averaged over 100 iterations. The blue line represents the end-state stability for Bayesian updating agents, while the gray line shows the end-state stability for groups with no updating. The dashed line shows the mean first-3-round stability for Bayesian updaters. Stability is measured as the proportion of pairwise dominance relationships which are consistent across all three rounds.

As seen in figures B8 and B9, while Bayesian updating produces some improvement in linearity and stability (compared to no updating, solid gray, or the first round, dotted red), particularly where initial self-assessment error is high, the effect is not nearly as dramatic as prior models, which tend to show strong linearity even in groups with identically sized individuals.

In our case, groups of identically sized individuals almost never form linear dominance hierarchies. We were able to generate this form of artificial linearity by setting opponent assessment error to be zero, self-assessment error to be very high, and wager bias, l, to be very high, such that after the initial, essentially random outcomes, agents arrived at high different estimates, but because l was so high, smaller-estimating individuals never invested enough to have a chance at winning and were thus trapped in their incorrect estimates. This is a very unstable condition, however, for example if some naïve, smaller agent were to stumble into this group, it would provide an opportunity for the original agents to improve their estimates and eventually all realize that they the same size as their group mates. This does highlight that under Bayesian updating, if agents' environment or contest strategy does not provide any opportunity to gain new, informative experiences, they can become trapped in estimates that might have been formed based on stochastic outcomes and very little information. This is an important feature of Bayesian updating, but the general point is that without the any feedback

on intrinsic ability, our assessment-only, Bayesian updating model has limited potential for hyper-stratification. It seems likely that most natural systems have some feedback mechanisms that allow winners to gain greater resources or social capital, which increases their true contest potential, in addition to increasing their estimates. Future models could incorporate estimate update and feedback on intrinsic ability to observe how these processes complement and interact.

Summary of Part 1

So far, we have explored the observations in Part 1 of the main-body—namely that 1) Bayesian updating produces winner-loser effects, 2) that Bayesian updating does not have a strong bias towards the loser-effect, that under Bayesian-updating, winner-loser effects 3) show a recency bias and 4) attenuate quickly in the face of new information, and that 5) Bayesian updating produces networks that are moderately stable, accurate, efficient, and linear. In all of these cases, these qualitative phenomenon were broadly observed in the parameter space, although it is important to note the edge-cases mentioned, and that winner effects generally disappear where effort is not a factor, or where contests are strongly dependent on chance. We now explore the sensitivity of the empirical predictions.

Part 2: Empirical Predictions

2.1 Size and Experience

As shown in figure 4a, we observe a discrepancy effect for winners in Bayesian updating, but not other updating paradigms. We repeated this experiment, this time including losers, but limiting 'treatment' opponents to either small and large agents (figure B10), iterating across the full range of parameters to observe whether this discrepancy effect always occurs (figure B11).

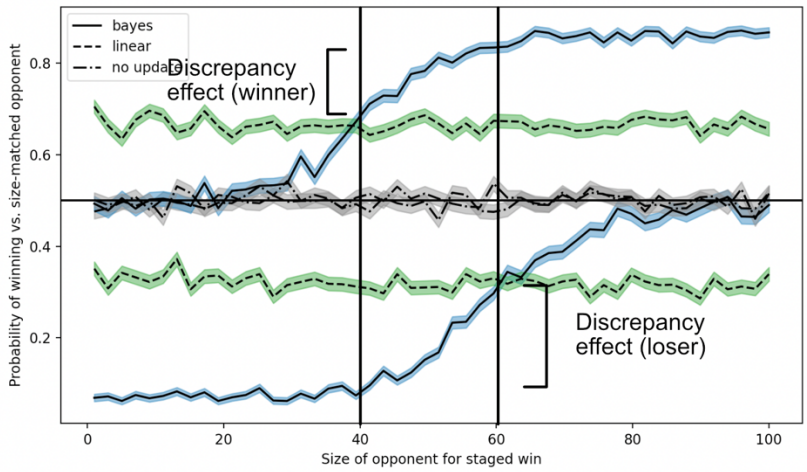

Figure B10: The discrepancy effect is measured as the difference between the winner/loser effect for a smaller 'treatment' opponent vs a larger 'treatment' opponent. These lines show the probability of winning against a second, size-matched 'assay' opponent, for three different updating strategies. This phenomenon requires that opponent size is factored into the strength of the winner-loser effect, as in the Bayesian updating approach in blue. Note that for both winners and losers, competing against a larger 'treatment' opponent results in a higher probability of winning the 'assay' contest.

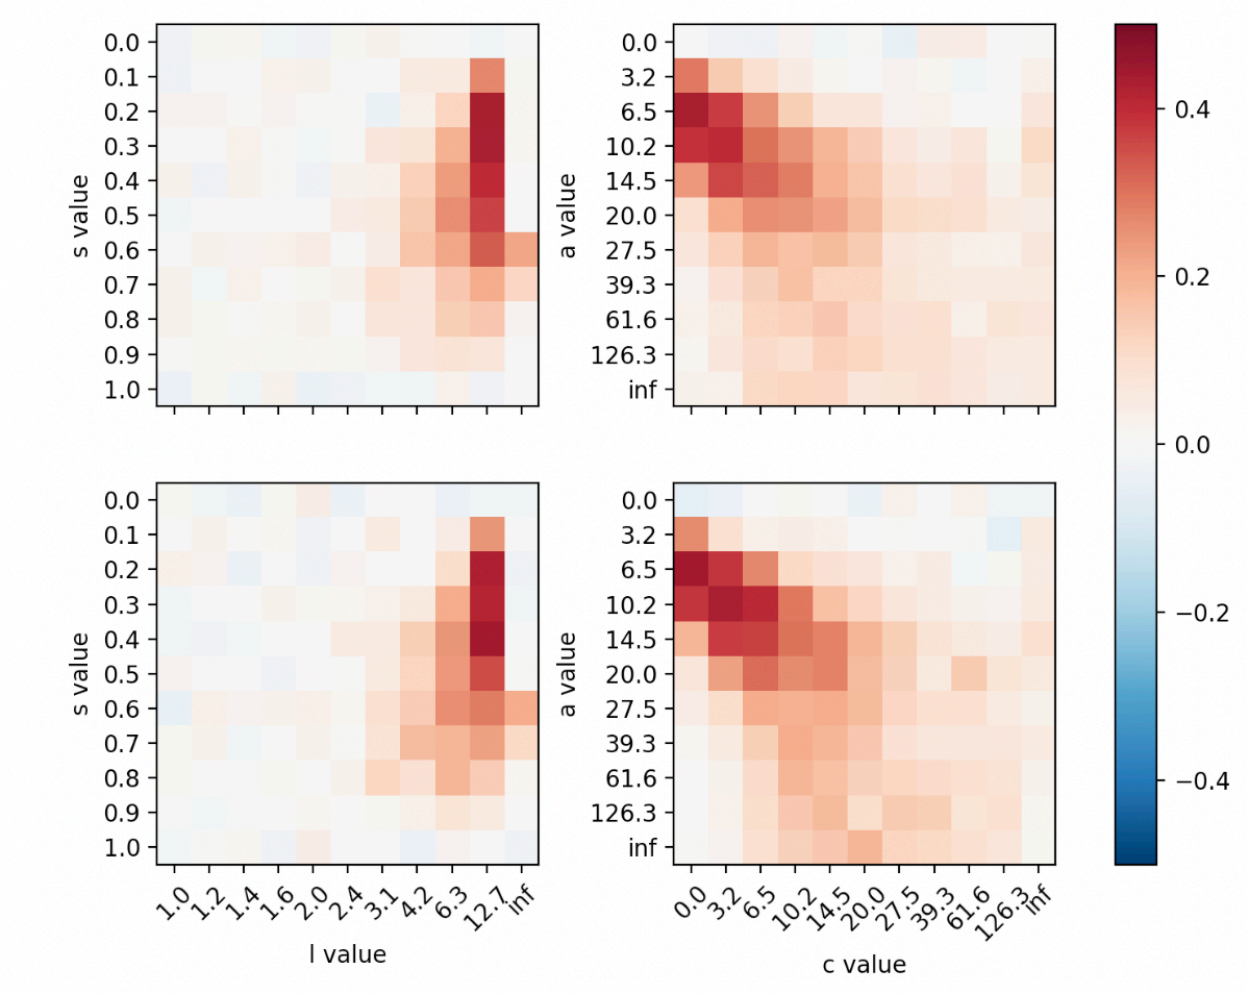

Figure B11: The discrepancy effect is broadly observed, but drops off quickly as l increases. Similar to the winner-loser effect in general (see Figure S5), the discrepancy effect is largely dependent on contests being strongly biased against the underdog ($l > 4$). Where the relationship between relative wager and winning is closer to linear, this effect is no longer observed.

Note that the discrepancy effect is only observed if both size and effort are factors, where self-assessment error is greater than 0, and opponent assessment error is lower than self-assessment error. This assessment limitation is because where opponent assessment is high, the opponent size-estimate in the size-matched 'assay' contest was largely random, meaning

1093 the focal agent effort was also very noise. This behavioral noise masks weak winner-loser

1094 effects.

1095

1096 We also show in figure 5b that winner effects are subject to prior experience, and the same is

1097 true for loser effects (figure B12). We can measure this "experience effect" by comparing naïve

1098 and experienced (in this case 10 previous contest) agents.

1099

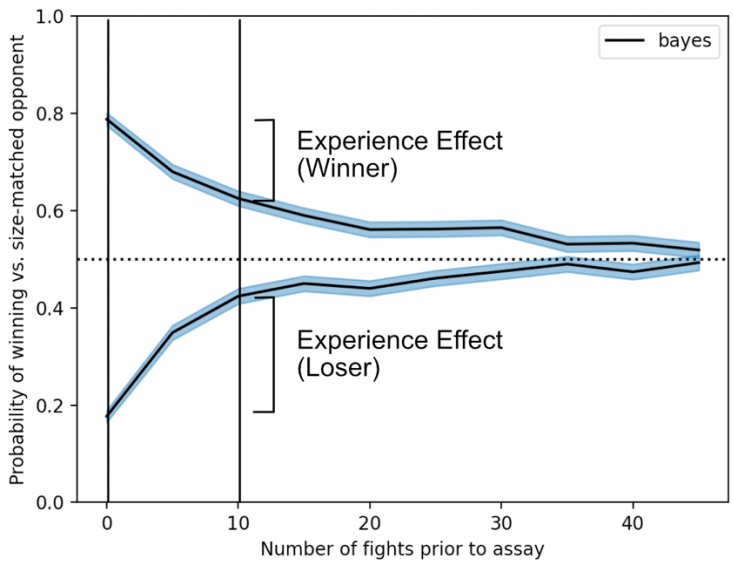

1100

1101 Figure B12: The experience effect is measured as the difference between the probability of winning a size-

1102 matched 'assay' contest after a 'treatment' contest for a naïve individual vs an experienced individual. In this

1103 case, the only difference is the number of contests agents have experienced prior to their 'treatment'

1104 contest. For both focal winners and focal losers, Bayesian updating results in a decreased winner/loser effect

1105 as experience (and thus prior confidence) increases.

1106

1107 As before, we repeated this simulation, averaging across 1000 iterations, comparing the naïve

1108 or experienced agents, across all contest parameters, to observe where this experience effect

occurred. This phenomenon is observable under roughly the same parameters as the discrepancy effect is observed, but with added noise introduced by the 10 pre-treatment opponents for the experienced individuals.

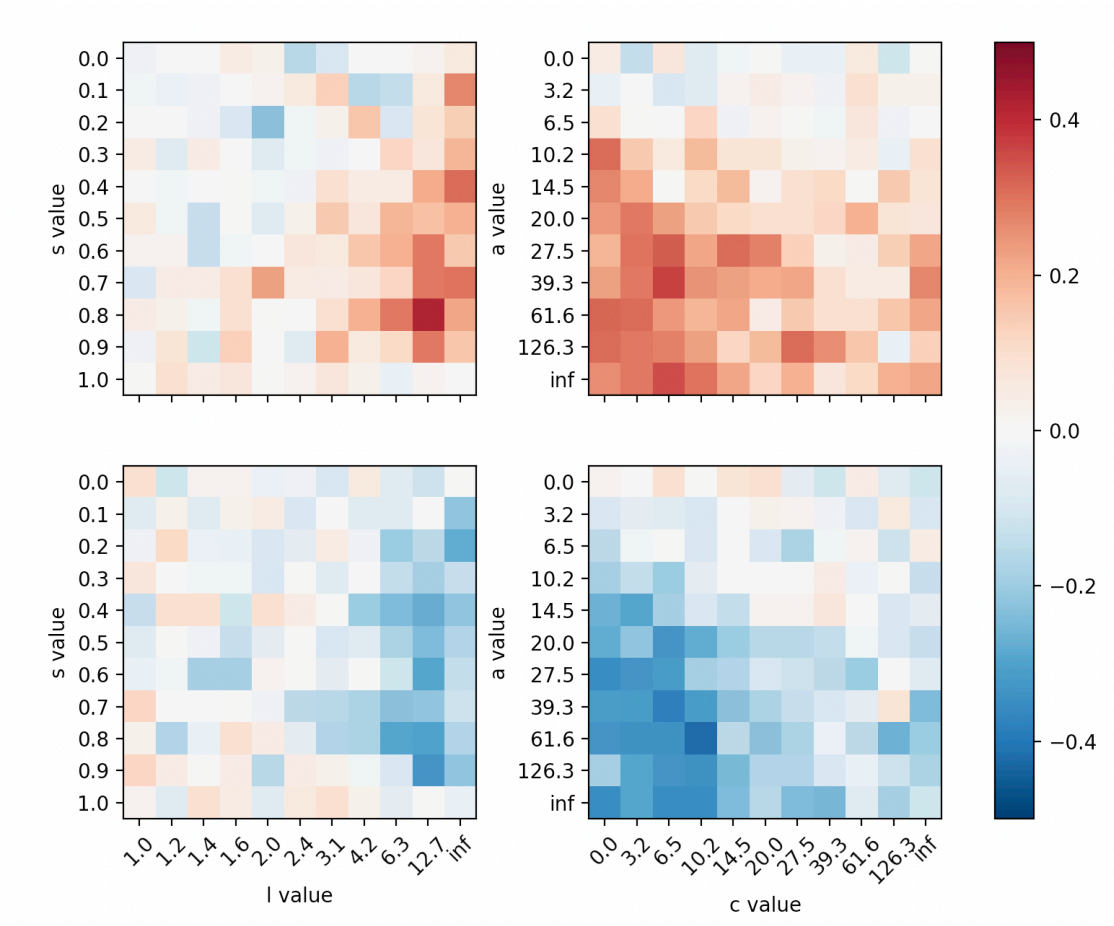

Figure B13: The experience effect is broadly observed wherever the winner-loser effect exists. The top row shows the strength of the experience effect (measured as the difference between naïve and experienced individuals) for winners over contest parameters s and l (left column) and assessment parameters σ_a and σ_c (right column), while the bottom row shows the same for focal losers.

Taken together, the discrepancy and experience effects are hallmarks of Bayesian updating, and indeed we would define any model which featured these two effects as "Bayesian-like", even if

the actual mechanism is not strictly Bayesian. We would expect these phenomena to occur in systems size and effort both contribute to contest outcome, and where naïve self-assessment error is roughly equal to or greater than error in opponent assessment.

2.2 Informed and uninformed priors

Finally, we explore under what conditions we should expect to see an effect between size and effort. We show in figure 5 that for agents using self-informed Bayesian priors, there is a strong correlation between size and effort. In the main body, we did not allow agents to update their self-assessment between contests, The results are qualitatively similar for the self-informed prior groups if updating occurs, although allowing Bayesian updating for all groups obviously introduces weak to moderate correlations between size and effort (figure S18).

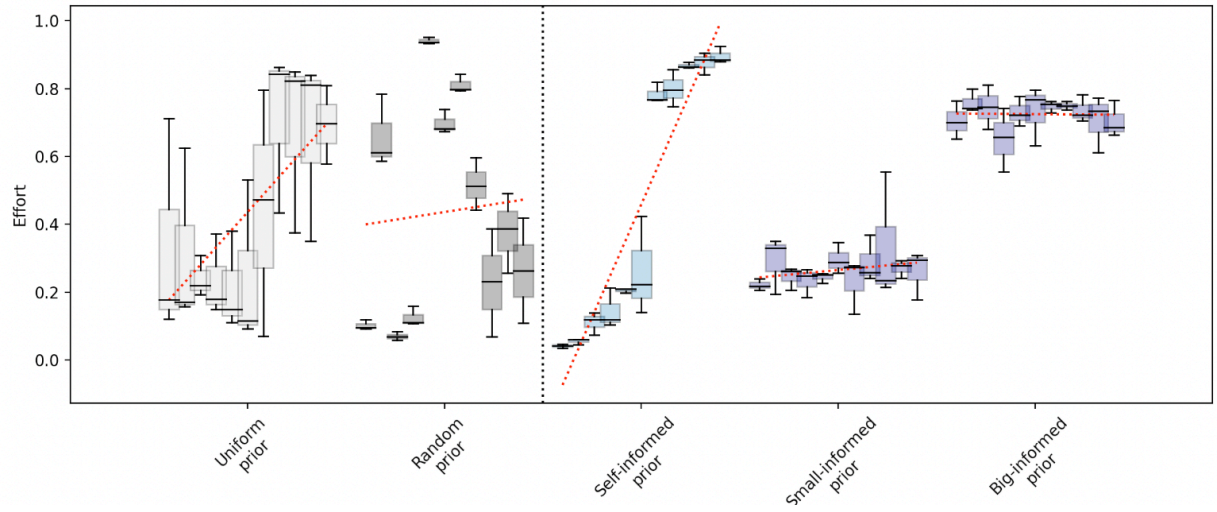

Figure B14: This modifies of the results figure 5 by allowing agents perform Bayesian updating after each repeated contest. Under these conditions, there are strong correlations for both uniform and self-informed priors, while random and externally-informed priors showed only weak correlations.

Focusing on the effect between effort and size observed in self-informed agents (and once again restricting updating), we repeat the simulation shown in figure 5 for every point in the parameter space, averaging the effect size across 1000 iterations (figure B15). We observe a positive effect in all conditions, except in unlikely edge cases (e.g., where the large animal wins without exception).

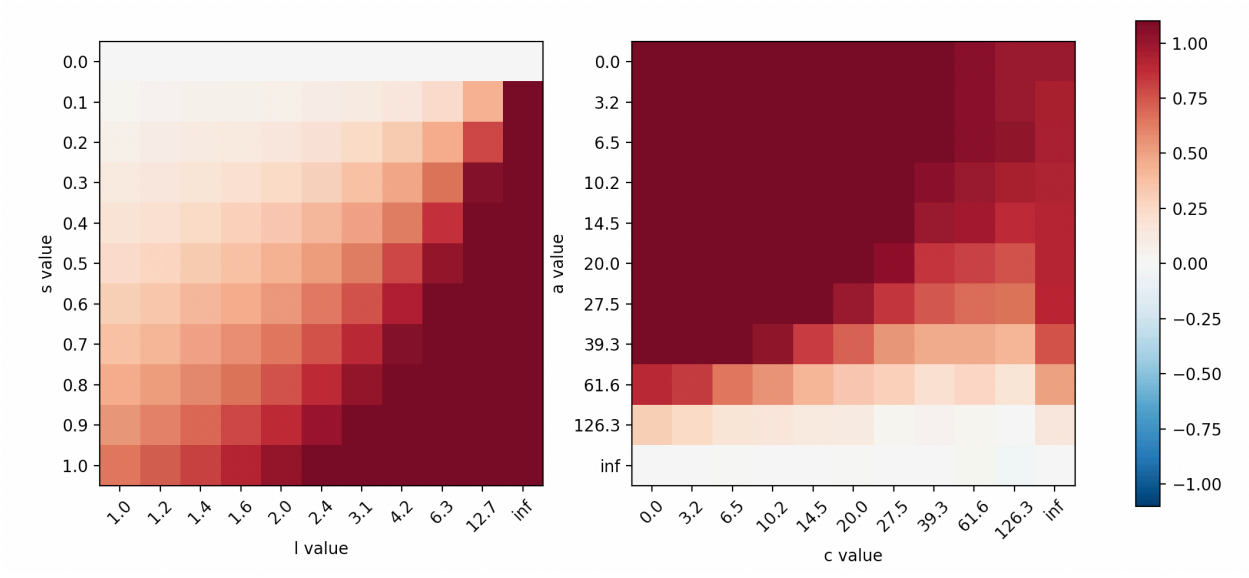

Figure B15: Size and effort correlates for self-informed agents across a range of parameter values. Here, the heatmap shows the strength of the effect (i.e., the slope of normalized size, $\frac{x}{100}$, and effort, simulated as in Figure 5, averaged across 1000 iterations. For utility, the heatmap is clipped at 1.1, such that any slope higher than that is normalized. Over all values where size contributes, there is a significant effect between size and effort.

It is also important to note that while this effect a consistent prediction of self-informed assessment, to measure this effect, it is necessary to isolate “effort” from possible confounds like energy budget or motivation. In our model, we typically think of *effort* as the amount of

time an individual animal stays in the contest, however this will often be linked to their available energy, which is frequently a function of individual size. This experiment would be most useful where *size* referred to something more specific, such as weapon size (as in figure A1), so that experimenters could control for intrinsic differences in the ability to persist in a fight. Alternatively, *effort* could measure the latency to engage in a contest, or some other non-energy-based metric. Obviously, applying theoretical predictions to empirical systems requires care and system-specific understanding to design experiments appropriately.

Summary

Bayesian updating predicts the existence of winner-loser effects, and several associated phenomena, however these predictions are most useful where individuals combine intrinsic ability and effort to determine the outcome, and where there is some (but not too much) uncertainty in own and opponent ability. These conditions likely apply to many social animals, and future empirical work should test these predictions in appropriate systems to observe whether natural behavior matches the predictions of Bayesian updating.

Appendix C – Parameters and Equations

Table 1: model parameters and abbreviations, reprinted here for reference

Symbol	Default	Range	Explanation
x : $(x, x_i, x_o, \bar{x}_i)$	Random	$1 \leq x \leq 100$	Size of an individual. x_i represents actual size for individual i, while x_o is the actual size of opponent. $\bar{x}$ denotes an estimated size
μ_x	50	$1 \leq x \leq 100$	The mean of the population from which individuals are drawn
σ_x	10	$0 \leq \sigma_x < \infty$	Standard deviation of population from which individuals are drawn. At $\sigma = \infty$, uniform distribution is used.
v : (v, v_i, v_o)	None	$0 \leq v \leq 1$	Effort, i.e., the amount of time/energy invested in a contest. As with size, subscripts distinguish effort by the focal individual v_i from their opponent's, v_o
w : (w, w, w)	None	$0 \leq w \leq 1$	Wager, or contest performance, a combination of the agent's size and selected effort. This could also be thought of as work, i.e., the amount of energy imposed on opponent.
s	0.7	$0 \leq s < 1$	Relative importance of size vs. effort, acts as a simple multiplier when summing size and effort.
l	6.3*	$0 \leq l < \infty$	Relative wager dampener, essentially the importance of luck, acting as an exponent on relative wager, such that as l increases, the probability of an upsets decreases
σ_a	20	$0 \leq \sigma_a < \infty$	*self-assessment error*, the accuracy of individuals' starting guesses for their self-assessment.
σ_c	3.1*	$0 \leq \sigma_c < \infty$	*opponent-assessment error*, the accuracy of individuals when guessing the size of opponents prior to contests
$P(\theta)$	-	-	Used to denote the prior, an estimated, discrete distribution of the probability of being some size.

Equations Sheet

Agent generation

$$x_i \sim N(\mu_x, \sigma_x) \tag{1}$$

Agent estimate generation

$$\bar{x}_i \sim N(x_i, \sigma_a) \tag{2a}$$

$$\mathrm{P_r}(\theta) = N(\bar{x}_i, \sigma_\theta) \tag{2b}$$

Effort strategy via opponent assessment

$$\bar{x}_o \sim N(x_o, \sigma_c) \tag{3a}$$

$$v_i = P(win\,|\bar{x}_i, \bar{x}_o) \tag{3b}$$

Calculating contest outcome as a function of relative wagers

$$w_i = s * x_i + (1 - s) * v_i \tag{4a}$$

$$P(upset) = \frac{1}{2} * \left(\frac{w_{min}}{w_{max}}\right)^l \tag{4b}$$

Linear updating strategy

$$\bar{x}_{i_{t+1}} = p\bar{x}_{i_t} + k \tag{5a}$$

$$X_{min} \leq \bar{x}_{i_{t+1}} \leq X_{max} \tag{5b}$$

$$p = \quad 0.1 * (X_{max} - \overline{x}_{i_t}), if\ win \tag{5c}$$

$$p = -0.1 * (\overline{x}_{i_t} - X_{min}), if\ loss$$

Bayesian updating of self-assessment

$$P_{post}(\theta|outcome) = \frac{P(outcome\mid\theta) * P_{prior}(\theta)}{P(outcome)} \tag{6a}$$

$$P(x_i = x|win) = \frac{P(win|x_i = x) * P(x_i = x)}{P(win)} \tag{6b}$$

$$P(win) = \sum_{x=0}^{x_{max}} P(win|x_i = x) * P(x_i = x) \tag{6c}$$

$$\bar{x}_i = E(\theta) = \sum_{x=1}^{x_{max}} x * P(x_i = x) \tag{6d}$$

Full effort strategy

$$v_i = \left(\frac{s\bar{x}_i + (1-s)*0.5}{s\bar{x}_o + (1-s)*0.5}\right)^l \qquad if\ \bar{x}_i \leq \bar{x}_o \tag{A1}$$

$$v_i = \left[1 - \left(\frac{s\bar{x}_o + (1-s)*0.5}{s\bar{x}_i + (1-s)*0.5}\right)^l\right] \qquad if\ \bar{x}_i > \bar{x}_o$$

Possible expansion of wager calculation

$$w_i = \beta_o x_i^a + \beta_1 v_i^b + \beta_2 (x_i * v_i)^c \tag{A2}$$

Full likelihood calculation

$$P(win|x_i = x) = P(upset), if\ w_i|\ x \leq w_o \tag{A3}$$

$$P(win|x_i = x) = 1 - P(upset), if\ w_i|\ x > w_o$$

Scaling functions

$$l = \text{cotangent}\left(\pi * \frac{(1-L)}{2}\right),\ \ 0 \leq L \leq 1 \tag{A4a}$$

$$\sigma_a = 20 \tan\left(\pi * \frac{A}{2}\right), 0 \leq A \leq 1 \tag{A4b}$$

Accuracy Calculation

$$\varepsilon = \sqrt{(\bar{x}_i + \sigma_\theta)^2 + (\bar{x}_i - \sigma_\theta)^2} \tag{B1}$$

References

Appleby, M. C. 1983. The probability of linearity in hierarchies. Animal Behaviour 31.

Arnott, G., and R. W. Elwood. 2008. Information gathering and decision making about resource value in animal contests. Animal Behaviour.

———. 2009. Assessment of fighting ability in animal contests. Animal Behaviour.

Bakker, T. C. M., E. F. Bruijn, and P. Sevenster. 1989. Asymmetrical Effects of Prior Winning and Losing on Dominance in Sticklebacks (Gasterosteus aculeatus). Ethology 82.

Beacham, J. L. 1988. The relative importance of body size and aggressive experience as determinants of dominance in pumpkinseed sunfish, Lepomis gibbosus. Animal Behaviour 36.

Becker, E. A., and C. A. Marler. 2015. Postcontest blockade of dopamine receptors inhibits development of the winner effect in the California Mouse (Peromyscus californicus). Behavioral Neuroscience 129.

Bierbach, D., M. Klein, V. Sassmannshausen, I. Schlupp, R. Riesch, J. Parzefall, and M. Plath. 2012. Divergent Evolution of Male Aggressive Behaviour: Another Reproductive Isolation Barrier in Extremophile Poeciliid Fishes? International Journal of Evolutionary Biology 2012.

Bonabeau, E., G. Theraulaz, and J. L. Deneubourg. 1996. Mathematical model of self-organizing hierarchies in animal societies. Bulletin of Mathematical Biology 58.

———. 1999. Dominance orders in animal societies: The self-organization hypothesis revisited. Bulletin of Mathematical Biology 61.

Briffa, M., L. U. Sneddon, and A. J. Wilson. 2015. Animal personality as a cause and consequence of contest behaviour. Biology Letters 11.

Chase, I. D., C. Bartolomeo, and L. A. Dugatkin. 1994. Aggressive interactions and inter-contest interval: How long do winners keep winning? Animal Behaviour 48.

Chase, I. D., D. Coelho, W. Lee, K. Mueller, and J. P. Curley. 2022. Networks never rest: An investigation of network evolution in three species of animals. Social Networks 68.

Chase, I. D., C. Tovey, D. Spangler-Martin, and M. Manfredonia. 2002. Individual differences versus social dynamics in the formation of animal dominance hierarchies. Proceedings of the National Academy of Sciences of the United States of America 99.

Cordero, M. I., and C. Sandi. 2007. Stress Amplifies Memory for Social Hierarchy. Frontiers in Neuroscience 1.

Dewsbury, D. A. 1988. Copulatory Behavior as Courtship Communication. Ethology.

Drews, C. 1993. The concept and definition of dominance in animal behaviour. Behaviour 125.

Dugatkin, L. A. 1997. Winner and loser effects and the structure of dominance hierarchies. Behavioral Ecology 8.

Ellis, L. 1995. Dominance and reproductive success among nonhuman animals: A cross-species comparison. Ethology and Sociobiology 16.

Enquist, M., and O. Leimar. 1983. Evolution of fighting behaviour: Decision rules and assessment of relative strength. Journal of Theoretical Biology 102.

Fawcett, T. W., and R. A. Johnstone. 2010. Learning your own strength: Winner and loser effects should change with age and experience. Proceedings of the Royal Society B: Biological Sciences 277.

Francis, R. C. 1990. Temperament in a Fish: A Longitudinal Study of the Development of Individual Differences in Aggression and Social Rank in the Midas Cichlid. Ethology 86.

Fuxjager, M. J., and C. A. Marler. 2010. How and why the winner effect forms: Influences of contest environment and species differences. Behavioral Ecology 21.
Fuxjager, M. J., J. L. Montgomery, and C. A. Marler. 2011*a*. Species differences in the winner effect disappear in response to post-victory testosterone manipulations. Proceedings of the Royal Society B: Biological Sciences 278.
Fuxjager, M. J., T. O. Oyegbile, and C. A. Marler. 2011*b*. Independent and additive contributions of postvictory testosterone and social experience to the development of the winner effect. Endocrinology 152.
Goetz, S. M. M., L. Tang, M. E. Thomason, M. P. Diamond, A. R. Hariri, and J. M. Carré. 2014. Testosterone rapidly increases neural reactivity to threat in healthy men: A novel two-step pharmacological challenge paradigm. Biological Psychiatry 76.
Goubault, M., and M. Decuignière. 2012. Previous experience and contest outcome: Winner effects persist in absence of evident loser effects in a parasitoid wasp. American Naturalist 180.
Graham, Z. A., and M. J. Angilletta. 2020. Claw size predicts dominance within and between invasive species of crayfish. Animal Behaviour 166.
Hemelrijk, C. K. 2000. Towards the integration of social dominance and spatial structure. Animal Behaviour 59.
———. 2002. Despotic societies, sexual attraction and the emergence of male “tolerance”: An agent-based model. Behaviour 139.
Hickey, J., and J. Davidsen. 2019. Self-organization and time-stability of social hierarchies. PLoS ONE 14.
Hock, K., and R. Huber. 2006. Modeling the acquisition of social rank in crayfish: Winner and loser effects and self-structuring properties. Behaviour 143.
Holekamp, K. E., and E. D. Strauss. 2016. Aggression and dominance: an interdisciplinary overview. Current Opinion in Behavioral Sciences.
Hsu, Y., R. L. Earley, and L. L. Wolf. 2006. Modulation of aggressive behaviour by fighting experience: Mechanisms and contest outcomes. Biological Reviews of the Cambridge Philosophical Society.
Hsu, Y., S. P. Lee, M. H. Chen, S. Y. Yang, and K. C. Cheng. 2008. Switching assessment strategy during a contest: fighting in killifish Kryptolebias marmoratus. Animal Behaviour 75.
Hu, S. W., Y. T. Yang, Y. Sun, Y. P. Zhan, and Y. Zhu. 2020. Serotonin Signals Overcome Loser Mentality in Drosophila. iScience 23.
Huang, S. P., S. Y. Yang, and Y. Hsu. 2011. Persistence of Winner and Loser Effects Depends on the Behaviour Measured. Ethology 117.
Ibuchi, K., and T. Nagayama. 2021. Opposing effects of dopamine on agonistic behaviour in crayfish. Journal of Experimental Biology 224.
Jackson, W. M., and R. L. Winnegrad. 1988. Linearity in dominance hierarchies: a second look at the individual attributes model. Animal Behaviour 36.
Johnsson, J. I. 1997. Individual recognition affects aggression and dominance relations in rainbow trout, Oncorhynchus mykiss. Ethology 103.
Kura, K., M. Broom, and A. Kandler. 2016. A Game-Theoretical Winner and Loser Model of Dominance Hierarchy Formation. Bulletin of Mathematical Biology 78.
Lan, Y. T., and Y. Hsu. 2011. Prior contest experience exerts a long-term influence on subsequent winner and loser effects. Frontiers in Zoology 8.

Landau, H. G. 1951. On dominance relations and the structure of animal societies: I. Effect of inherent characteristics. The Bulletin of Mathematical Biophysics 13.
Laskowski, K. L., M. Wolf, and D. Bierbach. 2016. The making of winners (And losers): How early dominance interactions determine adult social structure in a clonal fish. Proceedings of the Royal Society B: Biological Sciences 283.
Leimar, O. 2021. The evolution of social dominance through reinforcement learning. American Naturalist 197.
Li, C. Y., R. L. Earley, S. P. Huang, and Y. Hsu. 2014. Fighting experience alters brain androgen receptor expression dependent on testosterone status. Proceedings of the Royal Society B: Biological Sciences 281.
Lindquist, W. B., and I. D. Chase. 2009. Data-based analysis of winner-loser models of hierarchy formation in animals. Bulletin of Mathematical Biology.
Luttbeg, B. 1996. A comparative Bayes tactic for mate assessment and choice. Behavioral Ecology 7.
Maynard Smith, J. 1974. The theory of games and the evolution of animal conflicts. Journal of Theoretical Biology 47.
McDonald, A. L., N. W. Heimstra, and D. K. Damkot. 1968. Social modification of agonistic behaviour in fish. Animal Behaviour 16.
McNamara, J. M., R. F. Green, and O. Olsson. 2006. Bayes' theorem and its applications in animal behaviour. Oikos (Vol. 112).
McNamara, J. M., and O. Leimar. 2020. Game theory in biology: concepts and frontiers. Oxford University Press, Oxford, UK.
Mesterton-Gibbons, M. 1999. On the evolution of pure winner and loser effects: A game-theoretic model. Bulletin of Mathematical Biology 61.
Mesterton-Gibbons, M., Y. Dai, and M. Goubault. 2016. Modeling the evolution of winner and loser effects: A survey and prospectus. Mathematical Biosciences.
Momohara, Y., H. Minami, A. Kanai, and T. Nagayama. 2016. Role of cAMP signalling in winner and loser effects in crayfish agonistic encounters. European Journal of Neuroscience 44.
O'Connor, C. M., A. R. Reddon, I. Y. Ligocki, J. K. Hellmann, K. A. Garvy, S. E. Marsh-Rollo, I. M. Hamilton, et al. 2015. Motivation but not body size influences territorial contest dynamics in a wild cichlid fish. Animal Behaviour 107.
Okasha, S. 2013. The evolution of Bayesian updating. Philosophy of Science 80.
Olsson, O. 2006. Decisions under uncertainty: Bayesian foraging - A meeting held at Lund University in August 2003. Oikos (Vol. 112).
Parker, G. A. 1974. Assessment strategy and the evolution of animal conflicts. Journal of theoretical Biology 47:223–243.
Rowell, T. E. 1974. The concept of social dominance. Behavioral Biology 11.
Rutte, C., M. Taborsky, and M. W. G. Brinkhof. 2006. What sets the odds of winning and losing? Trends in Ecology and Evolution 21.
Schlinger, B. A., B. Palter, and G. V. Callard. 1987. A method to quantify aggressiveness in Japanese quail (Coturnix c. japonica). Physiology and Behavior 40.
Shizuka, D., and D. B. McDonald. 2015. The network motif architecture of dominance hierarchies. Journal of the Royal Society Interface 12.

Simons, N. D., V. Michopoulos, M. Wilson, L. B. Barreiro, and J. Tung. 2022. Agonism and grooming behaviour explain social status effects on physiology and gene regulation in rhesus macaques. Philosophical Transactions of the Royal Society B: Biological Sciences 377.

Smith, J. M., and G. A. Parker. 1976. The logic of asymmetric contests. Animal Behaviour 24.

Snyder-Mackler, N., J. R. Burger, L. Gaydosh, D. W. Belsky, G. A. Noppert, F. A. Campos, A. Bartolomucci, et al. 2020. Social determinants of health and survival in humans and other animals. Science.

Stevenson, P. A., H. A. Hofmann, K. Schoch, and K. Schildberger. 2000. The fight and flight responses of crickets depleted of biogenic amines. Journal of Neurobiology 43.

Strauss, E. D., A. R. Decasien, G. Galindo, E. A. Hobson, D. Shizuka, and J. P. Curley. 2022. DomArchive: a century of published dominance data. Philosophical Transactions of the Royal Society B: Biological Sciences 377.

Tibbetts, E. A., J. Pardo-Sanchez, and C. Weise. 2022. The establishment and maintenance of dominance hierarchies. Philosophical Transactions of the Royal Society B: Biological Sciences 377.

Van Doorn, G. S., G. M. Hengeveld, and F. J. Weissing. 2003. The evolution of social dominance I: Two-player models. Behaviour 140.

Whitehouse, M. E. A. 1997. Experience influences male-male contests in the spider Argyrodes antipodiana (Theridiidae: Araneae). Animal Behaviour 53.

Winberg, S., and P. O. Thörnqvist. 2016. Role of brain serotonin in modulating fish behavior. Current Zoology 62.