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 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, 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 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 resource holding potential (Chase et al. 2002; Tibbetts et al. 2022).

Winner and loser effects may be one of these social mechanisms determining the linearity and stability 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 similar, prior work has shown the loser effect tends to be stronger than the winner effect [reviewed in (Rutte et al. 2006)], and in some systems the winner effect has been shown to occur without an accompanying loser effect (and vice versa) [(McDonald et al. 1968; Goubault and Decuignière 2012) reviewed in (Hsu et al. 2006)]. In general, unless otherwise specified, we will discuss winner-loser effects together.

While winner and loser effects can be impacted by changes in individual contest potential 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 an individual's aggressive behavior [reviewed in Hsu, 2006]. This change in aggression suggests that individuals are modifying their self-assessment of their chances of winning in future contests based on the outcome of their current contest (Rutte et al. 2006). While some studies have linked changes in androgens and serotonin to winner-loser effects, the mechanisms behind post-fight shifts in behavior, as well as the broader evolutionary explanation of why changing behavior might be selected for, remain largely unknown. This fundamental gap in knowledge is further illustrated by the

conflicting empirical evidence on the duration of winner/loser effects. Winner-loser effects are typically observed as relatively short-term phenomena, lasting a limited time (i.e., minutes to days) (Chase et al. 1994; Johnsson 1997) or until some external disruption occurs (Stevenson et al. 2000). However, 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 understanding of the mechanisms explaining 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 focused on whether their existence was sufficient to explain the observed linearity of dominance hierarchies [reviewed in (Lindquist and Chase 2009)]. These models generally use a linear shift in estimate (and/or ability), in which after a contest, individuals' estimated (or actual) fight potentials increase by some fixed amount, (e.g., (Dugatkin 1997; Bonabeau et al. 1999)). This approach continues to be used in more recent models (e.g., (Hickey and Davidsen 2019)), though other models have developed more complex methods of contest assessment ((Hemelrijk 2002; Van Doorn et al. 2003; Hock and Huber 2006)). While these models are useful in exploring the impacts of winner-loser effects generally, they were not designed to identify the mechanisms of these effects, nor make specific predictions on animal behavior. Critically for mathematical ease, these models have typically conflated self-assessment and contest ability. However, empirical data show that winners and losers typically change their aggressiveness in subsequent fights suggesting that self-assessment alone may be sufficient to describe winner-loser effects, and we wanted a model which could isolate the effect of self-assessment. Second, some models have relied on complex learning algorithms for the agents, requiring perfect knowledge of opponent history. We sought to find a simpler and perhaps more biologically plausible mechanism to explain the persistence of winner-loser effects.

Bayesian updating provides the optimal modeling framework to explain winner/loser effects: it depends solely on changes in individual assessment and existing empirical research suggests that many animals use Bayesian processes to make decisions in other contexts [[McNamara et al, 2006, Okasha 2013, Olsson 2006, Luttbeg 1996). In brief, Bayesian updating provides a precise mathematical framework to modify an existing estimate of the state of the world based on new information (McNamara, 2006). 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. This is surprising, because Bayesian updating can account for both the *how* and the *why* of winner-loser effects: Under Bayesian updating, winner-loser effects are an expected side-effect of the central process of animals assessing and establishing their social status. 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, to better determine how much to invest in future contests. The purposes of this model were to first, explore the features and limitations of Bayesian updating as it relates to winner-loser effects, to show whether such a framework is plausible. To do so, we test whether it recreates known empirical features of dominance hierarchies, which tend to be linear and stable, as well as the recency biases of winner/loser effects and the higher relative strength of loser effects as shown in the empirical literature. Our second goal was to identify testable predictions that could empirically distinguish Bayesian updating from alternative mechanisms of dominance establishment and winner-loser effects. Here, we focus on measurable attributes of fights, such as fight intensity and duration and compare these measures across staged contests where the agents use either Bayesian updating or other updating mechanisms like linear shifts in self-assessment [as in (Dugatkin 1997) (Bonabeau et al. 1999), or 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.

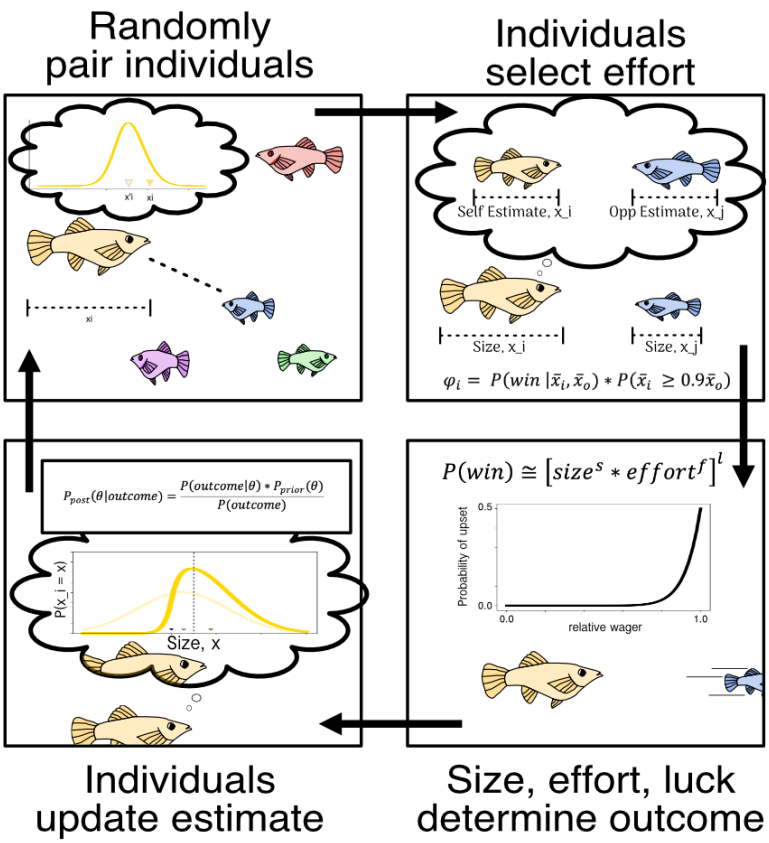

Figure 1: Graphical overview of conceptual framework. ***Top left****, Individuals with fixed size (unknown to self) and an estimated self-assessment are paired randomly.* ***Top right****, each individual estimates their own and opponent's size to determine their effort, based on their estimated probability of winning.* ***Bottom right****, the size and effort of each agent determines*

their relative wagers, based on the relative size and effort of the smaller individual. 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). The outcome is determined by whether a randomly generated value is below this probability. ***Bottom left****, based on this probabilistic outcome, agents update their self-assessment, multiplying their prior assessment by the calculated likelihood of the observed outcome, summed across all possible sizes. This process is repeated many times for all possible pairings to allow dominance hierarchies to emerge.*

General approach

We envision a scenario where a group of animals compete for dominance in war of attrition contests. In our case, we imagine a group of fish of varying size, but individual agents could represent any species (mice, humans, etc.) competing in a fixed social group where there is uncertainty around one's own and their competitor's contest abilities. 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, which we call their 'estimates' (Fig 1, top-left). 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 (Fig 1, top-right). Third, contest outcomes are determined probabilistically based the relative size and effort of each participant (Fig 1, 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 (Fig 1, bottom-left). This process is then repeated. Each step of the model is described in detail below.

Table of Abbreviations:

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 the time/energy invested in a contest, scaled by the size of the agent. This could also be thought of as work, i.e., the amount of energy imposed on opponent.
s	1	$0 \leq s < \infty$	Relative size dampener, acts as an exponent such that as s increases, so does the importance of relative size

f	1	$0 \leq f < \infty$	Effort dampener. Acts as an exponent, such that as f increases, so does marginal benefit of increased effort relative to the opponent
l	13	$0 \leq l < \infty$	Relative wager dampener, essentially the importance of luck, acting as an exponent on relative wager, such that as l increases, the probability of upsets decreases
σ_a	10	$0 \leq \sigma_a < \infty$	*awareness*, the accuracy of individuals' starting guesses for their self-assessment.
σ_c	10	$0 \leq \sigma_c < \infty$	*acuity*, 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.
Φ	-	-	Standard notation to denote the cumulative distribution function used to calculate confidence
p	1.1	$0 \leq x \leq 100$	Proportional shift after a contest, used in linear updating
k	0	$0 \leq x \leq 100$	Constant shift after a contest, used in linear updating
r	10	$1 \leq x < \infty$	Proportional shift after a contest, used in linear updating

Generation of agents

In order to model the development of self-assessment for each iteration, we first randomly generated agents with some intrinsic win potential, called *size*, drawn from a normal distribution, such that

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

Each agent had an estimate of their own size, $\bar{x}$, formed by drawing from a normal distribution centered around their actual size:

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

where σ_a reflects individual *awareness*, i.e., the precision with which they can estimate their own size. Individuals then calculated a starting prior as a normal distribution centered on their estimate, with standard deviation, $\sigma_\theta = \sigma_a$,

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

In addition to their estimate of size, all agents had an effort strategy and an estimate updating strategy, which we explain below. Groups of *n* simulated agents were paired across r rounds of

(n-1) contests, such that in every round, every agent faced every other agent, for a total of (n-1 $*$ r) contests, with the order of contests being shuffled for each round.

Effort strategy

The purpose of our model was to explore updating strategies, but in order to do this, we needed to establish how agents would determine their *effort* in a contest, v_i, representing the amount of time/energy an animal spends on a contest before yielding. Empirical observations [Hsu et al, 2008] and theory [Maynard Smith & Parker, 1976; Enquist & Leimar, 1983] suggest that individual 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 estimated the size of their opponent, by drawing their estimate of opponent size, $\bar{x}_o$, from a normal distribution centered on their opponent's true size, and standard deviation reflecting their *acuity,* σ_c*, i.e., the precision of that assessment,*

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

Agents then set their effort, v_i, based on their estimated probability of winning, such that

$$v_i = P(win\,|\bar{x}_i, \bar{x}_o) * P(x_i \geq 0.9x_o) \tag{3b}$$

Here effort, v_i, is equal to the conditional probability of winning—given an agent's estimates about their own and their opponent's sizes—multiplied by their estimated probability that their size is within 10% of their opponent's size (calculated using the confidence of these respective distributions, see Eq. 3d)

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,

$$v_i = \left(\frac{\bar{x}_i}{\bar{x}_o}\right)^{sl} * P(\bar{x}_i \geq 0.9\bar{x}_o), \qquad if\ \bar{x}_i \leq \bar{x}_o \tag{3c}$$

$$v_i = \left[1 - \left(\frac{\bar{x}_o}{\bar{x}_i}\right)^{sl}\right] * P(\bar{x}_i \geq 0.9\bar{x}_o), \qquad if\ \bar{x}_i > \bar{x}_o$$

where, s and l are simulation parameters used in several equations to account for the influence of *relative size* and *luck* respectively, on contest outcomes. Here, *luck* means how close the two agents' respective sizes and efforts need to be for there to be a reasonable probability of upset. Although the actual probability of winning depends on both relative size and individual effort (see contest calculation below), with this effort strategy, agents approximate the probability of winning by assuming that both they and their opponent will invest the same amount in the contest.

The second term in Eq. 3b and 3c, $P(\bar{x}_i \geq 0.9\bar{x}_o)$, accounts for the confidence of an individual's estimates, calculated as the probability that the individual is within at least 10% of the opponent size using a cumulative distribution function of the combined normal distributions:

$$P(x_i \geq 0.9x_o) = \Phi_{\mu_\delta,\sigma_\delta}(-0.1 * x_o) \tag{3d}$$

Because an agent's estimates of both their own size and opponent size are represented by normal distributions, these two distributions can be combined to calculate a new normal distribution representing the probability density of the difference in their sizes, where the new mean, μ_δ, represents the estimated difference in size $(\bar{x}_i - \bar{x}_o)$, and σ_δ is the sum of the standard deviations of the component distributions, i.e., *acuity* and *awareness*, $(\sigma_c + \sigma_a)$. In Equation 3d above, to calculate the probability that x_i is at least 0.9 times the size of x_o, agents use the cumulative distribution function, i.e., the sum of the probability over a set range of their size estimate distribution.

Determining contest outcome

For a given contest, the winner was decided probabilistically (Figure 1, bottom right), based on the relative size and effort of each agent, combined to define each agent's *wager*, w_i:

$$w_i = \left(\frac{x_i}{x_{max}}\right)^s * v_i{}^f \tag{4a}$$

Wager, w_i, and *effort,* v_i as terms are quite 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 relative size, and the parameters s and f allow us to control the relative importance of *size* and *effort*, respectively. Because these exponents always operate on fractions between 0 and 1, we can conceptualize them as "bias" parameters, where, as they increase, larger inputs (of size or effort) are exponentially more successful than smaller inputs (Supplemental Figure S1a provides further intuition on the behavior of these parameters).

Because in both our model and natural observations, contests are determined by when the subordinate individual yields, and contest success if a function of relative, not absolute, ability, it was 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 a third 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. In this way the agent with greater relative size and effort (termed the *favorite)* usually won, but the agent with the smaller wager (the underdog) could win, particularly if the two wagers were close. It is also important to note that the smaller wager was not necessarily the smaller sized individual; if a relatively smaller agent contributed far higher effort, they could have the larger wager, w_{max} , and thus be the *favorite* for the purposes of our calculations.

This approach to settling contests allowed us to vary the relative contributions of size, effort, and stochasticity in order to observe how this changed the behavior of the group (i.e., how much individuals invested in contests, whether dominance relationships were stable, etc.). For most purposes, we set the values of these bias parameters at s = $5, e = 1, l = 13$ in order to match empirical observations (Beacham 1988; Bierbach et al. 2012), that upsets were rare, and contests were usually predictable based on size (Supplemental Figure S1a).

Updating self-assessment post contest

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

Fixed Estimate (No updating)

The simplest possible update strategy is to not alter your initial estimate. We modeled this scenario, where we kept 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 model estimate updating, (see, e.g., (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. To make this more comparable to our Bayesian updating approach, we modified this approach to set limits on the maximum and minimum possible size (X_{min}, X_{max}), such that their new estimate, $\bar{x}_{i_{t+1}}$, was a function of their previous estimate, shifted by some factor,

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

where

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

The scaling parameters are theoretically unbounded, but to more closely compare to the dynamic shifts of Bayesian updating, in practice we set $k = 0$, and

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

such that individuals shifted their estimates of their own size as a function of their distance from the maximum or minimum possible size. After updating the mean estimate, we generated a new normal distribution, centered around that estimate, based on their *awareness*, σ_a, as in Eq. 2b.

Bayesian updating

In contrast with simple heuristics like linear updating, 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 described above as a discrete probability distribution across all possible sizes. 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. Thus, for our purposes, 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. Equation 6a means that to update the probability distribution of size, at each size x, we multiply the previously estimated probability of being size x by the likelihood of winning the previous contest at that size. The probabilities of winning, for every possible size, are then divided by the sum of the probabilities of winning for all size values (i.e., the total probability of winning). Thus, 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 Eq. 4b. Because probability of winning is a function of the relative of the smaller wager, the probability of winning goes from being equal to the probability of upset, to its inverse, once winning is conditioned on having the bigger wager.

$$\begin{aligned} P(win|x_i = x) &= P(upset), if\ w_i|\ x \leq w_o \\ P(win|x_i = x) &= 1 - P(upset), if\ w_i|\ x > w_o \end{aligned} \tag{6c}$$

Note that there is an inflection point 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 (see Equation S2).

The prior and likelihood are then divided by the total $P(win)$, which in our case is the sum of prior $*$ likelihood for all possible sizes,

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

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

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

Where E is the maximum likelihood estimate of the size distribution. Agents then calculate an estimated confidence, σ_θ,

$$\sigma_\theta = \sum_{x=1}^{x_{max}} P(x_i = x) * (x - \bar{x}_i)^2 \tag{6f}$$

These values are then used to compute their effort in the next contest (using Equation 3c, with the updated σ_θ taking the place of initial awareness, σ_a).

Total parameter space

Taken together, our possible parameter space comprises the contest outcome parameters (s, f, l), reflecting the influence of *size*, *effort*, *luck*; the assessment parameters (σ_a, σ_c), representing *acuity* and *awareness*; and, for linear updating, the value of the *scalar* and *proportional* parameters (k, p). Even with our simplifying assumptions, this represents a somewhat daunting parameter space, especially since these parameters are theoretically unbounded. To aid in exploring the parameter space, for the contest and assessment parameters, which all varied between [0, ∞), we used a scaling function (Supplemental Equation S1) so that we could input parameters as values between either [1,-1] for (s, f, l)—such that the inflection point at e.g., $s = 1$ was maintained at the scaled value of 0—and [0,1] for (σ_a, σ_c). By doing this, it was more straightforward to explore the full range of possible parameter values.

<u>Model Analysis</u>

The goal of our model was 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?

Our core approach to answering these questions was to run many (e.g., 1000, but see results) iterations of social simulations while extracting behavioral metrics that would be tractable in and relevant to an empirical context, such as contest intensity. Because these are simulations, we were also able to directly inspect the hidden variables (e.g., an agent's actual estimate of

own size) that were driving the behavior. We organize the metrics according to the level of complexity at which they are observed.

Agent metrics

For each agent, we recorded true intrinsic ability (referred to simply as *size*), as well as their self-assessment—the estimated probability distribution of their size—at every timepoint t, with the calculated mean and variance of the distribution.

Contest metrics

For every contest, we record the outcome, as well as the intensity of the contest, determined by the lower effort, $min(v_i, v_o)$, of the two competitors. We also recorded the effort, v, and wager, w of each agent for each contest.

Hierarchy metrics

For the full history of social interactions, we track the linearity, stability, accuracy, and efficiency of dominance contest interactions. In order to provide network metrics throughout the simulation, we binned the history of interactions into 3-round sub-histories, using a sliding window, and calculated metrics both within bins, and (where applicable) across the full history. *Linearity* is calculated using Appleby's formula (Appleby 1983), which calculates the proportion of triads that are transitive. For *stability*, we first calculate binary dominance-subordinance status for each pairwise relationship, then calculate the proportion of time bins for which the binary dominance network, calculated within that time bin, matches the binary dominance network calculated for the entire simulation. For *accuracy*, we calculate ordered dominance, using an ELO ranking (Albers and De Vries 2001; Neumann et al. 2011), then calculate the spearman correlation between actual size and dominance rank. For our metric of *efficiency*, we simply measure the mean intensity of contests within each time bin, where here a time bin was one round of all pairwise contests.

Implementation

All code was written in python. The random number generator for contest outcome was from the *numpy* library. For ELO ranking we used an *elosports* module (pypi.org/project/elosports). All code and results are available on github (github.com/aperkes/BayesBattleBots).

Results

Part 1: Matching Empirical Observations

1.1 Bayesian updating predicts that the loser effect is stronger than the winner effect.

We first tested whether agents using Bayesian updating experienced winner-loser effects at all. To do so, we modeled our simulation on empirical approaches by staging a win or loss between a focal agent and a "treatment" opponent of known size (in this case, the same size as the focal individual). We then measured whether there was a winner/loser effect by pairing these focal winners and losers against new "assay" opponents (again size-matched to the focal individual). We could then quantify the strength of the winner/loser effect as the proportion of focal agents

which won against their second size-matched assay opponent (where no winner-loser effect would yield 50% wins). Here we found obvious winner-loser effects, in that prior winners were more likely to win subsequent contests, while prior losers were more likely to lose (Fig 2). Additionally, consistent with the strong empirical support documenting that the effects of losing a contest tend to be stronger than the effects of winning a contest, in our model, the loser effect is significantly stronger (Fig 2). In our model, this phenomenon is driven by two mechanisms, stemming from the empirically based assumption that contests are determined by the relative size of the smaller competitor, such that differences in the perceived *relative* sizes of winners or losers vs their opponents (e.g., 70% vs 50% of the larger agent) will result in disproportionate outcomes, even if the *absolute* perceived size differences (50:75 vs 25:50) are the same (see Supplemental Figure S2).

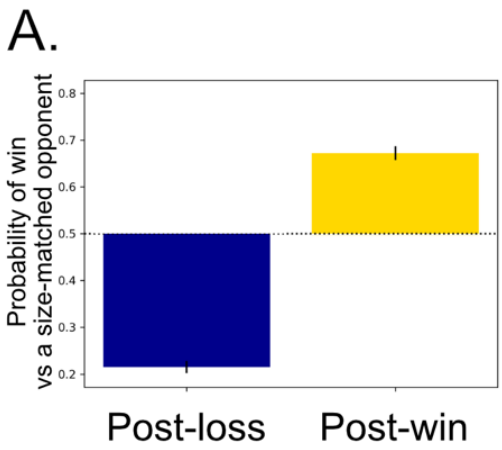

Figure 2. *A. Loser effects are stronger than winner effects. This plot shows the proportion of focal agents (averaged from groups of n=1000) winning against a naïve size-matched fish, after experienced either a loss (blue) or a win (gold). The dotted line shows the expected proportion of wins (0.5) if this second contest were unbiased. In this context, loser effects are stronger than winner effects, measured as the distance from 0.5.*

<u>1.2 Bayesian updating produces a recency bias in winner-loser effects</u>

Although winner-loser effects vary across species, there is strong empirical evidence that many species show a recency bias with the outcomes of more recent contests having a stronger effect than earlier contests on current behavior. We found that our model recreated these empirical observations. Again, emulating empirical tests of the winner-loser effect, we simulated contests between naive, size-matched agents and then paired the subsequent winners and losers against new, naïve, size-matched individuals. To simplify this simulation, all individual priors were centered on their true size, and agents correctly assessed opponent size, regardless of their acuity, σ_c. Focal agent then faced new, size-matched, naive individuals, where each contest generated a new branch in the focal agent's path. This process was repeated 6 times,

for 1000 individuals to explore all possible outcomes of the sequence of 6 contests. This time, we measured winner-loser effects directly as the agents' own size estimates. We observed that winner-loser effects have an experience-dependent recency bias (Figure 3), such that after first winning a contest and then losing a contest (WL), agents updated their size estimate to be lower than if they had lost and then won (LW). This recency effect was based on the fact that the initial losers (LW) won while investing less in their second contest than the initial winners (WL); the strength of this effect depends on the specific parameters (see Supplemental S.3), but under all parameters, losing then winning resulted in estimates that are greater than or equal to winning than losing.

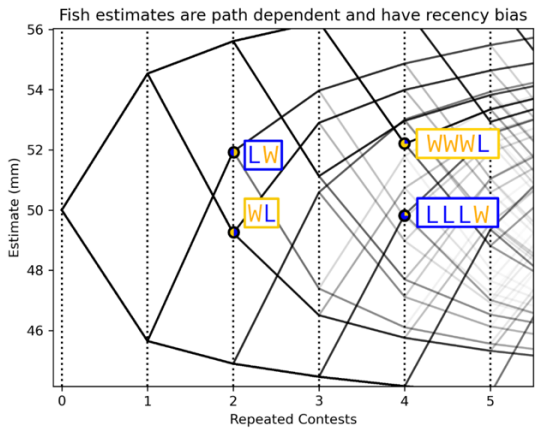

Figure 3: Divergent experience leads to stochastic formation of self-assessment. Agent self-assessments are plotted as they compete in repeated contests against naive, size-matched opponents. Each vertical dotted line represents a contest, with the branching paths showing the estimates following a win or loss. The boxes note specific experiences (e.g., LW: a loss followed by a win), highlighting the recency effect, and how the recency effect diminishes with increased experience. The box outlines are colored to represent the agent's experience, prior to the penultimate contest. Here, fight parameters are [S,E,L] = [0,0.7,-0.8]

<u>1.3 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 provides a mechanism through which to reconcile these apparent discrepancies based on the size of the initial opponent, the number of new events agents experience following this "treatment" contest, and the method used to assay the strength of the winner effect.

We again simulated a forced-win experimental approach, this time allowing a simulated focal agent of size $x_o = 50$ to win a staged contest against an opponent of variable size ($x_o = (40,50, or\ 60)$. As before, we measured the winner-effect directly, as the agent's own self-assessment. Following this initial "treatment" contest, focal agents were allowed to interact

with additional naïve individuals to observe how their estimate changed over time. In this context, it is apparent that the size of the opponent has a strong effect both on both the degree of the initial shift in self-assessment and how long this boosted estimate persists (Figure 4A), with additional experiences causing agents' estimates to tend towards their actual size. Exactly how long the winner effect lasts above some given value is influenced by the nature of these post-fight experiences (Supplemental Figure S3), but even after 100 contests, our "distant past" winners' self-assessment was higher than their actual size, and those who won against bigger opponents had significantly higher self-assessments than those who won over smaller ones.

To test whether these differences in self-assessment would be detectable in experimental contexts, we simulated experiments of varying sample sizes, in which we measured the performance of these focal winners and losers (i.e., the proportion of fights won) against either size-matched or smaller opponents. In this context, the size of the assay opponent was critical (Figure 4B). The "distant-past" winners and losers could be reliably distinguished by comparing the proportion who won in contests against size-matched individuals. This was true even for small sample sizes (e.g., 5 winners vs 5 losers). In comparison, if focal agents were tested against smaller opponents, it was extremely rare to capture any significant difference for groups smaller than 100.

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 was 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 specific methods used influence whether winner-loser effects are observed.

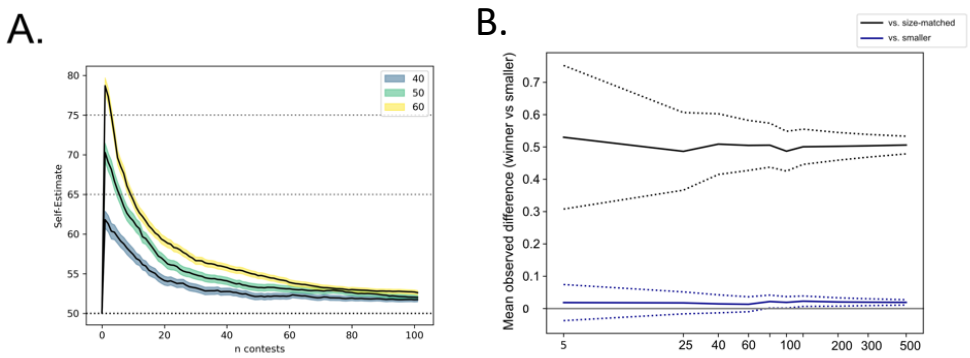

Figure 4: Strength and persistence of the winner effect is a function of social information and experimental power. A. This plot shows focal agents' estimates following a staged win against agents of varying size. The shift in the focal agent's estimate following a win is greater when their opponent was larger. B. This plot shows the mean estimated winner effect in 1000 simulated experiments, calculated by measuring the difference in win rate for focal winners and losers against a size-matched agent (black) or a smaller agent (blue), with the dotted lines showing the standard deviation of estimates.

1.4 Bayesian updating produces linear, stable, and efficient dominance networks

Winner-loser effects are important not just for individual behavior but also the formation of dominance hierarchies. We found that our simulated agents using Bayesian updating recreated well-known empirical observations patterns of dominance networks. Matching observations of typical natural dominance hierarchies, we found that our simulated agents recreated stable linear hierarchies across a range of parameters tested (Supplemental Fig S4). This tendency toward stabilization was sufficiently strong that groups of agents formed linear dominance hierarchies even in the absence of size differences (Figure 5b, yellow line). There were certain edge cases where other outcomes occurred (see Supplemental S4) but linearity was a general phenomenon across the parameter space. We additionally found that the linearity of hierarchies of our simulated agents rapidly increased (Fig 5b) while the intensity of contests decreased (measured as the lower of the two efforts, v_{min}) (Fig 5c). Low intensity contests are a hallmark of many empirical observations as individuals often seek to avoid serious injury, and in our simulation, these occurred as the smaller agent in each contest began to recognize that they were unlikely to win and thus yielded quickly.

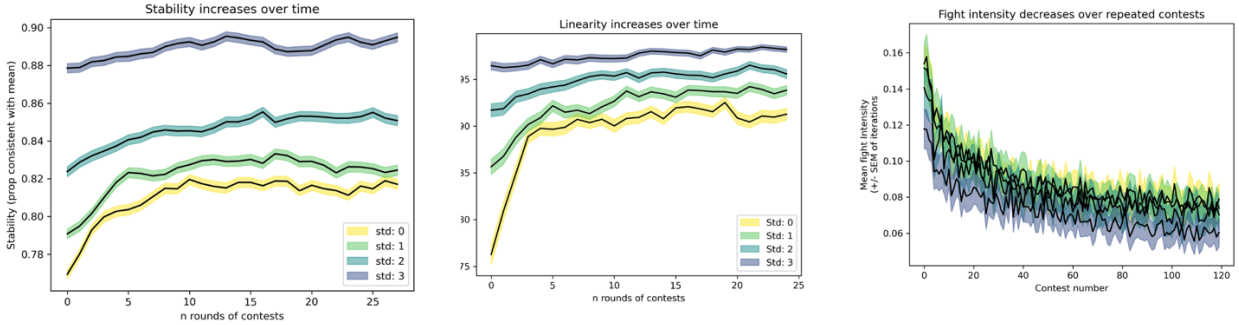

Fig 5: Emergence of stable dominance hierarchy through Bayesian updating. ***A. Network stability increases over time.*** *Each line represents the mean stability across 1000 simulations for a given amount of variation in population size (with yellow, $\sigma_x = 0$, where all individuals are equally sized). Stability is measured by binning the simulation into 3-round windows and comparing the consistency of dominance rank across bins (calculated using an ELO score). Stability tended to increase, particularly where real size differences were low.* ***B. Linearity of dominance networks increase over time.*** *As with stability, the linearity, measured as the proportion of triads in the dominance network that are transitive, following Appleby's approach.* ***C. 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. Although there was relatively little difference in contest intensity, based on size variance, more variable populations did have lower intensity contests, which is to be expected, since assessment is easier when the absolute difference is greater.*

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 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 investigated 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 potential updating mechanisms: Bayesian, linear updating, or fixed shifts (no updating). To test how these different updating mechanisms might be experimentally distinguished, we simulated contests of focal individuals against treatment opponents of the full range of possible sizes and then measured the strength of winner-effects using a size-matched assay opponent match. In this context, we found that the effect of opponent-size on the strength of the winner effect provides a method to distinguish the different potential updating mechanisms: agents using Bayesian updating show increases in winner effects when their initial opponent is larger, whereas these effects are not seen if the agents use linear updating or fixed shift (Figure 6a). This is in line with the discrepancy effect, which is a well-known characteristic of Bayesian updating processes whereby more surprising (and therefore informative) outcomes lead to a more dramatic shift when updating the estimate, whereas linear updating has no mechanism to incorporate opponent size.

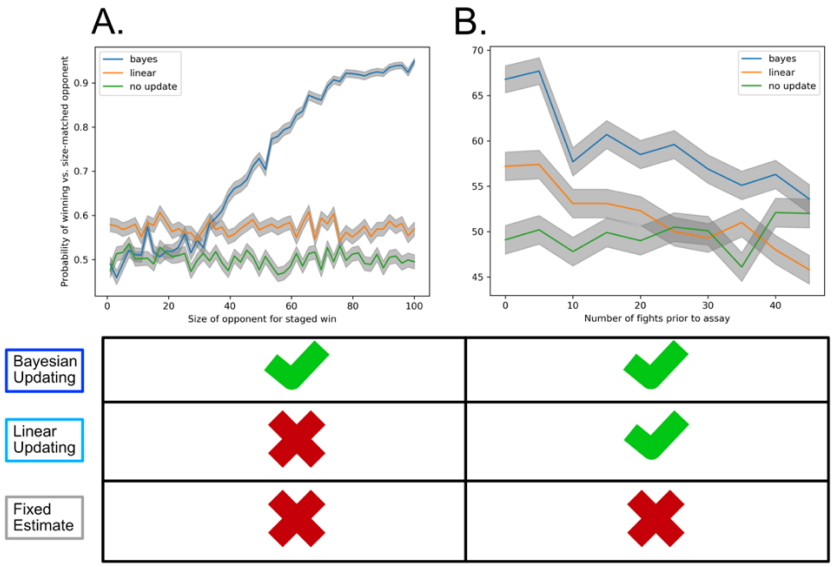

Figure 6: Simulated experiments show that updating strategy can be distinguished by the effect of opponent size and focal agent experience. A. Under Bayesian updating, the strength of the

winner effect (the proportion of agents winning a size-matched contest after a staged win), is a function of opponent. 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. C. The inset table shows the predictions for each updating. approach, with green checks denoting a predicted significant effect of treatment on the strength of the winner effect.

We next explored how the updating mechanism used by agents alters the influence of past experience (i.e., the number of prior contests) on winner-loser effects. To do so, we allowed simulated focal agents to experience between 0 and 50 fair contests against naïve opponents, in which the probability of upset determined by size and effort (as described in Equation 4b), thus creating variation in the extent of each focal agent's experience. We then assessed the strength of the winner-loser effect by allowing this focal individual to fight against a size-matched individual in a final test match. In this context, agents using Bayesian updating show weaker winner effects with greater experience (i.e., more contests, Figure 6b, dotted line). This is consistent with the 'confidence' effect of Bayesian updating: more experienced individuals are expected to have more confident priors and thus be less influenced by new information. To our surprise, this same experience effect was also evident in individuals that were using linear updating schemes (Figure 6b, dashed line), which appears to be driven by the fact that under linear updating, estimates tend to diverge over time, being either extremely large or small (see Supplemental Figure S9). Thus, while an experience effect is a core prediction of Bayesian updating, 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 discrepancy effect (i.e., opponent-size) is the most straightforward single test for Bayesian updating.

<u>2.2 Informed prior estimates drive consistent individual and context dependent differences</u>

Often when behaviorists test for whether animals are using Bayesian strategies, they are asking specifically whether (and how) animals form prior estimates of the general state of the world and then apply that to a new context, for example the quality of a novel potential mate or some new foraging patch. In this sense, the question is not about how individuals update their prior estimate in the moment, but whether animals form prior estimates from past experience and use them in new contexts. With this in mind, we simulated experiments that would allow us to distinguish whether individuals had informed priors, independent of their updating strategy after a contest, and what information these priors were based on.

To do this, we first simulated agents that used one of four sources to form their prior own-size estimate: 1. An uninformed prior, the uniform distribution; 2. a mis-informed prior, a normal distribution centered on a uniformly random variable; 3. A self-informed prior, a normal distribution where mu approximated the agents size (See Equation 2a & 2b), or a peer-informed prior, a distribution that reflected the distribution of nearby agents (which we set to be either all large or all small). Of these, (3) and (4) both represent "informed" priors.

We then conducted simulated experiments in which we measured 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 was assayed 3 times, each time against a different, naïve size-matched opponent (drawn from the same prior method group). Because we had 3 trials for each agent, we could measure both persistent individual differences and context-dependent variation in the amount of time agents spent before yielding (i.e., *effort*). We allowed agents to update their estimates after each contest, but results were not dependent on the updating strategy.

In these simulated experiments, we found key differences in aggressive behavior which successfully characterized groups based on the source of their prior (Figure 7). Specifically, in groups of individuals using either self-informed or mis-informed priors, there were consistent individual differences in contest effort (i.e., significant variation within the group based on individual). Peer-informed agents could show individual differences if these agents were drawn from different peer groups (small peers vs large peers) but did not show differences when drawn from the same peer group (small peers vs small peers). In self-informed agents only, the individual variation in effort correlated with own-size, while in peer-informed agents (drawn from different groups) variation in effort correlated with peer size. Thus, by controlling for peer group and characterizing individual size and effort, it is possible to identify whether individuals form prior estimates and, if so, what information they use to inform these estimates.

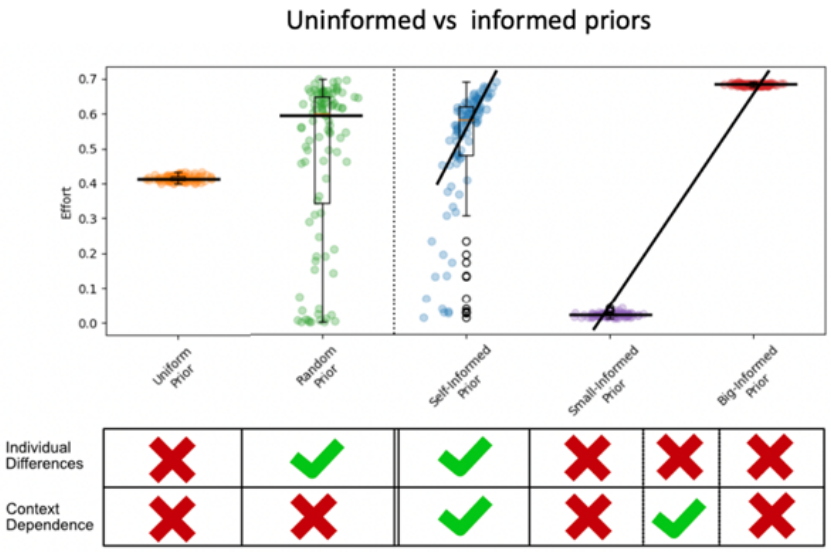

Figure 7: Agent behavior reveals the information agents used to form prior Top: Contest intensity measured for 5 different simulated groups, each using a different starting prior. The two left-most populations have "uninformed" prior, either with all individuals having a uniform prior, or each individual having their own randomly selected prior. In the right three groups, individuals have "informed" prior, using either internal information (central cluster, blue) or external information (two rightmost clusters). Black lines show the best fit line between effort and size for each population. The inset table shows, for each group, which effects are predicted to be significant for contest effort (green check).

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 in natural social systems. 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 model broadly matches the known features of dominance interactions and winner-loser effects. Empirical observations vary with regard to the existence of winner-loser effects, how long those effects last when they are found, as well as the relative strength of winning or losing on future contest outcomes, but our model shows that the properties of Bayesian updating provide a mechanism that can reconcile the disparate results seen in the empirical literature. We show that the strength and duration of winner/loser effects follow directly from the information given to the individual before, during, and after the contest. This provides an important framework for interpreting existing winner-loser effect results, as well as planning future experiments.

In our model, varying the amount of information given to the individual is enough to match the strength and persistence of winner/loser effects to any given values, but it is likely there are additional mechanisms that could be added to our model to better match empirical observations of social dynamics. One obvious expansion would be to allow individual ability (*size*) to vary over short and long timescales. In our model, we deliberately kept agent size fixed in order to focus on the importance of their self-assessment, but in nature, contests use energy and often yield new resources and/or territory, all of which alter an individual's short- and long-term contest potential. Physiologically, there are also multiple cascades of "fight or flight" neuro-endocrine responses that are activated during a contest and in response to wins and losses, and the metabolic state of the individual likely changes independent of any permanent change in self-assessment. These processes add short- and long-term uncertainty as to the state of an individual and thus the probability of winning a contest. All of this could be incorporated into our Bayesian framework, by allowing *size* to vary over short and long timescales in response to contests and allowing agents to maintain both "current" and "average" estimates of their self-assessed contest ability which are updated according to new information. This would allow us to explore how this added uncertainty changes the value of Bayesian updating.

A second important addition would be to allow individuals to gather more information from contests between peers. While our model generated social networks that function similar to those observed in nature, our agents were only maintaining estimates about their own size and did not attempt to "remember" any information about opponents. Empirical work has shown that in most animal societies, individuals are keenly aware of the social interactions happening around them, and that allowing individuals to observe contests reinforces linearity and produces dynamics that cannot be experimentally replicated with the "round robin" tournaments on which our model is based (Lindquist and Chase 2009). While social

eavesdropping and individual recognition increase the complexity of the model, these could easily be incorporated as additional Bayesian priors which agents update and incorporate into their decisions, and doing so would make it possible to make predictions for how access to both intrinsic and social information, as well as individual recognition, changes the behavior of social networks.

An implicit assumption in most discussions of “the” winner-loser effect is that there is a single “effect” of winning or losing; however, it is far more likely there are multiple effects that occur across different timescales. The existence of Bayesian updating for self-assessment does not rule out other concurrent mechanisms, like rapid changes in metabolism due to fight-flight responses, however Bayesian updating does provide an intuitive and powerful explanation for a wide range of variation observed in animal behavior, which can be scaled to include multiple processes that occur over both short and long timespans. Given that Bayesian updating is the most accurate possible approach to improve an estimate in the face of uncertainty, Bayesian updating can serve as useful null hypothesis for any repeated behavior, allowing us to make both general and specific predictions for how individuals and groups behave over time, as we have shown here with our predictions that past experience or contest information (i.e., intensity) should impact the strength and persistence of winner-loser effects. To our knowledge, there is no existing empirical work testing how these factors impact behavior in social contests, providing several possible directions for future experiments. By developing and testing Bayesian models for social behavior, we can better 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

S.1 Contest parameter intuition

As described in the methods are settled by based on relative wager (Eq) where wager is a function of relative size and effort, each raised to an exponent between 0 and infinity. In order to explore the parameter space more evenly, we use a scaling function

$$l = \cot\left(\frac{(L+1)}{2} * \frac{\pi}{2}\right), -1 \leq L \leq 1 \qquad \text{(S1)}$$

This effectively maps $L = 1$ to $l = 0$, L $= 0$ to $l = 1$, and $L = -1$ to $l = \infty$, as shown in Figure S1a. We can think of these parameters as magnifying or flattening differences in the different aspects of wager, which varies the importance of these aspects. To help illustrate this, we simulated a contest between two fish, one larger ($x_i = 60$) but low effort ($v_i = 0.3$), another smaller ($x_j = 50$) but high effort ($v_j = 0.6$). When s is large, the larger agent tends to win, when f is large, the higher investing agent tends to win, and as l increases, the bias against the "underdog" increases.

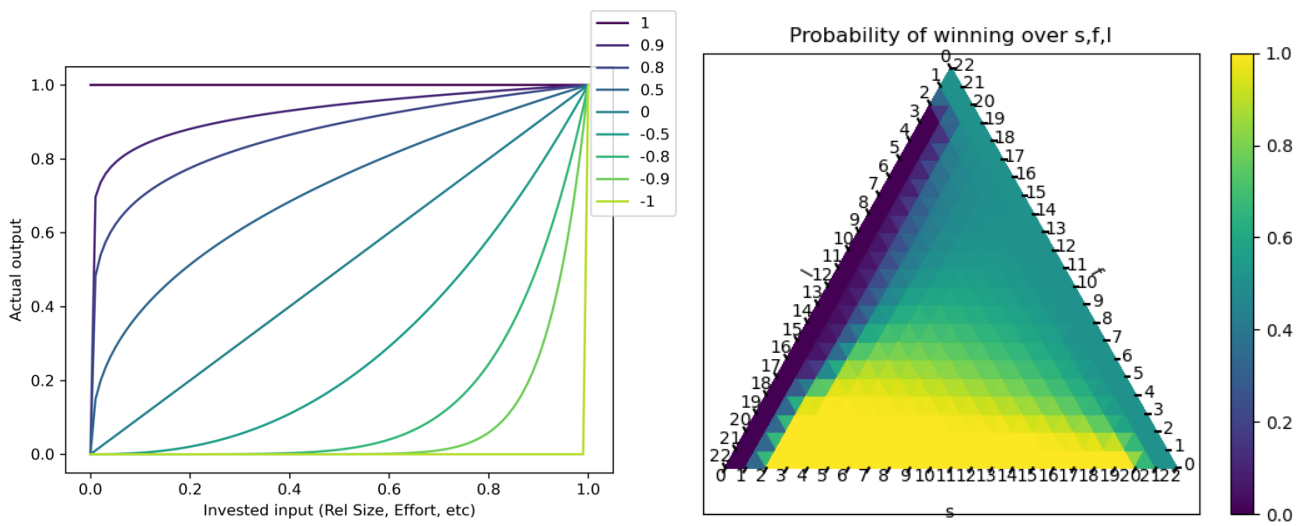

Supplemental Figure S1: Contest parameters. Panel A shows how relative size, effort, and relative wagers are scaled based on the value of s,f or l, respectively. Here again we use Equation S1 to scale parameters, such that 1 represents s^0 while -1 represents s^∞. Panel B shows the probability of a larger fish (size=60, effort=0.3) winning against a smaller, smaller competitor investing greater effort (size=50,effort = 0.6). Axes move evenly from -1 to 1 in the scaled parameter space (Supplemental Equation S1).

S.2 Understanding why the loser effect is stronger

To illustrate why the loser effect is stronger, we will first examine this simplified scenario where, following a size-matched contest, 50cm individuals either increase or decrease their estimate of their own size by 25cm, based on winning or losing, respectively. For simplicity in exploring the effect of effort, we assume here that the probability of winning, $P(win)$ is strictly equal to relative effort, $\frac{v_{min}}{v_{max}}$. As Figure 4b illustrates, after the staged contest, the perceived probability of winning a new contest against a naïve, size-matched individual is 25% for the loser and 67% for the winner (The probability for a naïve agent would be 50%). Here, assuming their effort is simply equal to their perceived relative size, we see that the loser decreases their effort by 0.25, while the winner only increases their effort by 0.17, thus the loser will have a relatively lower wager ($w_{rel} = 0.5$) against a naïve opponent, than that same opponent would have against a winner ($w_{rel} = 0.\overline{6}$). In other words, equal (but opposite) shifts in absolute size estimates lead to unequal shifts in perceived relative size, leading to different relative effort, and thus different relative outcomes, for winners versus losers.

This relative effort effect is sufficient to drive differences in the winner and loser effect, but Bayesian updating further increases the relative strength of the loser effect due to asymmetry in likelihood functions across the range of relative sizes. As a reminder, when calculating likelihood functions, agents calculate their probability of winning against an opponent of fixed size, x_o, for all possible values of their own size, x_i. Here, to highlight size effects, we assume likelihood of winning is strictly equal to relative size. First, as Figure 4c shows, likelihood functions are asymmetrical with respect to the size of the opponent for winning and losing. This asymmetry occurs because, once an agent is larger than the size opponent, as the agent's assumed size increases, the relative size of the smaller opponent, and thus the likelihood of winning, decreases hyperbolically. Simplifying equation 6c, we have

$$P(win\,|x_i = x) = 1 - \frac{x_o}{x_i},\ for\ x \geq x_o \qquad \text{(S2)}$$

In the extreme cases of the smallest or biggest possible individual (in the range $[1,100]$) competing against a mean individual, the size of the minimum individual is 2% of the mean ($\mu_x = 50$), whereas a mean individual is 50% of the maximum sized agent. This flattens the right side of the likelihood function (Figure 4c). Since the new, posterior size estimate is found by multiplying the likelihood by the prior, this asymmetry in the likelihood function means, all else equal, posterior estimates decrease following a loss by more than they increase following a comparable win. Taken together, the two effects of relative size and relative effort combine to make the loser effect stronger than the winner effect, consistent with empirical observations.

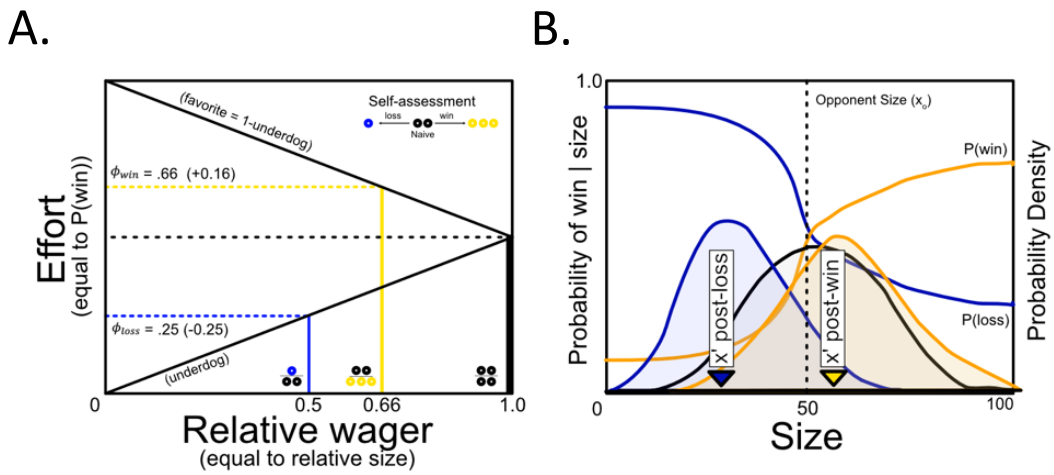

Supplemental Figure S2 A. The effects of relative scaling are such that an equal decrease in estimate leads to a disproportionate decrease in effort. B. The flattening of the likelihood function at larger sizes means losses are more impactful.

S.3 Understanding the Recency Effect

While it is consistent with empirical findings, it may seem counterintuitive mathematically that our Bayesian updating model shows a recency effect, since Bayesian updating is typically order invariant. This recency effect results from the fact that effort, v_i, is determined by an agent's estimate of their own size, such that the agents that won their first contest, but lost their second (WL) invested more heavily in their second contest (after having won their first) and thus the upset was even more surprising. In Bayesian terms, prior winners (WL) experienced a steeper likelihood function in their second-contest loss than the LW agents did in their first contest, and thus WL agents decreased their prior more dramatically. In other words, in a Bayesian updating paradigm, the recency bias comes not from any particular temporal bias, but from a difference in the experience itself, vis-à-vis the perceived likelihood of winning given the relative effort given by the focal agent.

It is important to note that our model lacks any explicit measure of time, only the number of contests experienced. Thus, with no time-bias, winner and loser effects remain fixed in our model until some new experience alters them. In this context, the strength of the recency bias decreases with each additional contest agents had experienced previously, making winner-loser effects weaker as experience increases (see Figure 6b below). Indeed, there are some conditions (e.g., WWWL vs LLLW, highlighted in Figure 3) where the observed recency bias is not sufficient to overcome the dramatically different prior experiences and so no winner or loser effect is observed, in the sense that LLLW agents are still less likely to win a size-matched contest (as compared to a naïve individual).

S.4 Understanding persistence of winner effect

In the main results we discuss the impact of opponent size, but Bayesian updating also allows us to model individual estimate confidence (i.e., the narrowness of the distribution of possible size estimates). We tested the impact of estimate confidence by repeating the above experiment, but instead of varying the size of the focal opponent, we varied the precision of individuals initial estimates, σ_θ, while keeping the accuracy (i.e., *awareness*, σ_a) the same (See Equation 2a,2b). Here, not surprisingly, the confidence of the prior at the time of the treatment contest attenuated the winner effect but had relatively little impact on the final state.

We also explored the how changing the post-contest environment changes the persistence of the winner effect. In any empirical observation of the winner-loser effect, individuals are typically placed in an initial contest (in which they win or lose) and then assayed again in a subsequent winner-loser test match. However, in between these two contests, individuals are still experiencing their environment and any additional information or experiences in this environment may influence the individual in their subsequent contest. In the empirical world, relevant information come from any number of sources; in our model, we explored how access to other interactions prior to the subsequent winner-loser test match alters the strength and persistence of the measured effects. We created focal winners of size $x_i = 50$, and allowed these focal individuals to win against treatment opponents (this time of size $x_o = 60$), but this time, we placed our focal agents with groups of varying numbers of naïve opponent agents (i.e., n = 3,5,10,20, or 40). Here, the sizes of naïve agents were sampled randomly from a normal distribution, $N(\mu_x = 50, \sigma_x = 10)$. This allowed all focal agents to compete against several other naïve agents, thus gaining new information and updating their estimates with each contest. [DID NAÏVE INDIVIDUALS UPDATE THEIR ESTIMATES?] In this simulated experiment, the number, and body size of naive individuals our focal agents faced over the course of 100 contests had a strong effect on how much and how quickly the winner-effect from the initial focal contest was attenuated. Briefly, focal agents that interacted with naïve agents with greater variation in their body sizes more rapidly improved their estimates of their own sizes thus attenuating winner effects more quickly than focal individuals who only interacted with similarly sized opponents (supplemental). (Figure 5b).

B.

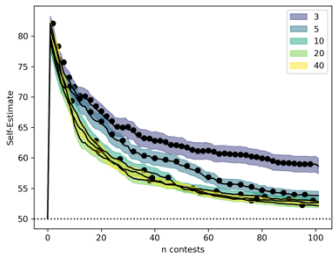

Supplemental Figure S3. Here, the focal winner is placed in groups of varying size. The x axis shows the number of total contests, during which focal agents face all other opponents in their group to form their round, and then face all opponents again and again (each round in a different, randomized order) until they reach 100 contests. The circles along the lines mark the

start of each round of opponents. The persistence of this shift depends on the variation in the actual sizes of the other agents in the group.

S.5 Linearity persists or increases across parameters

In general, we found that our simulated agents recreated stable linear hierarchies across the range of parameters tested (Supplemental Figure S2). However, there were certain edge cases where other outcomes occurred: for example, if the error of agents' starting estimates, σ_a, was high and/or the size differences between individuals were small (Figure 2a), there were more frequent upsets (both based on size and based on relative wager) resulting in less-stable hierarchies, especially in early rounds. Where size differences were large (Figure 2a,2b), or if the size-bias parameter, s, in Eq. 4a was high, such that any size difference could entirely determine outcome (i.e., s approaches infinity and l is set sufficiently low that outputs are not extremely stochastic), then hierarchies were entirely deterministic based on size and thus generally stable from the start. This initial stability is inconsistent with most empirical dominance hierarchies, which tend to feature an initial period of high intensity, lower-consistency contests, during which dominance relationships are established (Tibbetts et al. 2022), followed by stabilization. In our simulation, under most parameter configurations, the tendency toward stabilization was sufficiently strong that groups of agents formed linear dominance hierarchies even in the absence of size differences (Figure 2b, yellow line). This was especially true where small differences in effort could strongly determine outcomes (i.e., a high effort-bias parameter, f, in Eq. 4a), and in these contexts we observed the strongest contrast between the initially chaotic network (with less linearity, higher mean contest intensity and less consistent winners/losers) and the eventual stable state, (Supplemental Figure S3c)

To test whether the stabilization effect that we observed based on Bayesian updating was consistent across parameters, we simulated 100 dominance hierarchies across the range of possible parameters. Using the scaling function above, we tested every combination from -1 to 1 in intervals of 0.1. For any dominance hierarchy, the random expectation of linearity is 0.75 (i.e., 3 of 4 triads should be linear just by chance). Dominance hierarchies tended to begin more linear when size was a larger component (Figure S2a). We allowed these groups to interact across 20 rounds, and again measured linearity (Figure S2b). By taking the difference between this initial linearity, and the end-state linearity, we observed that in general, linearity either increased or stayed the same (Figure S2c) This increase was most dramatic where size had little to no effect on contest outcome, and where effort was a strong factor. Dominance networks that started out strictly linear obviously did not increase.

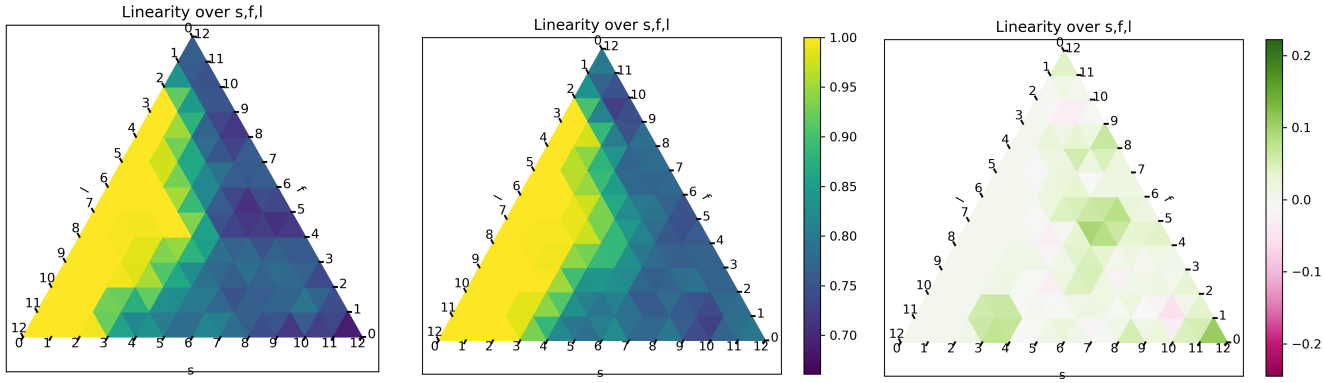

Supplemental Figure S4: Linearity. These plots show the results of simulating dominance hierarchies across various parameters. For each parameter set, we repeated the simulation 100 times. Panel A shows the linearity in the first three rounds of contests. Panel B shows the linearity after 20 rounds. Panel C shows the difference.

S.6 How is effort determined?

Agents invested in a contest based on their perceived probability of winning, which was calculated according to (Equation 3c), as their probability of winning given their estimate, multiplied by the (rough) probability of their estimate being correct. To demonstrate this, we simulated a contest between a smaller agent and a bigger agent, with various awareness and acuity parameters (Figure S3). As an agent's awareness decreased (i.e., the confidence of its self-assessment) so did effort. As acuity increased, effort decreased, since the smaller agents correctly identified that their probability of winning was low.

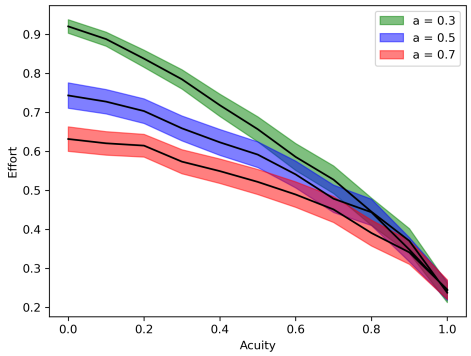

Supplemental Figure S5: The effects of acuity and awareness on effort. This plot shows the mean effort (from 1000 simulated contests) of a smaller agent ($x_i = 40$) against a bigger agent ($x_j = 60$) as a function of acuity. The different lines represent agent of different awareness.

Because they change both the magnitude and variance of effort, acuity and awareness also had an impact on the linearity, and the stabilization of dominance networks. As before, we simulated many groups across the range of acuity and awareness parameters (Figure S4), measuring linearity before (S4a) and after (S4b) repeated interactions. Here the most dramatic stabilization was observed when acuity was low and there was moderate awareness (S4c). Where agents had strong self-assessment, dominance networks were linear from the start, high acuity tended to increase stabilization, since it allows accurate self-assessments to be used effectively.

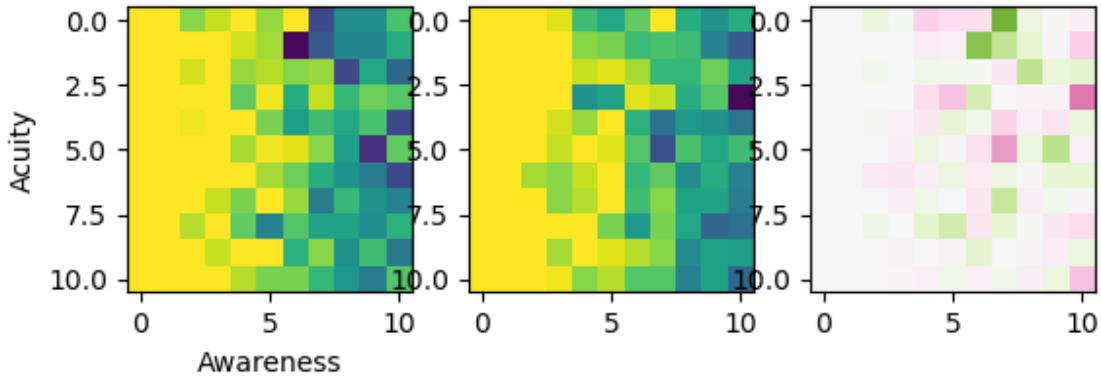

Supplemental Figure S6: Linearity across possible parameters acuity and awareness. As in Figure S2, these panels show the measured linearity before (A) and after (B) 20 rounds of contests. Here, the axes move exponentially from 0 to infinity, and represent the standard deviation around the mean when assessing opponents. Acuity 10 (i.e., $c = \infty$) is the same as pure self-assessment, since agents simply assume the opponent is average size.

To confirm that our stabilization effect was not an artifact of mutual assessment, we repeated our linearity measures with moderate awareness but minimal acuity. Under these conditions, agents still increased their linearity over time (Figure S5).

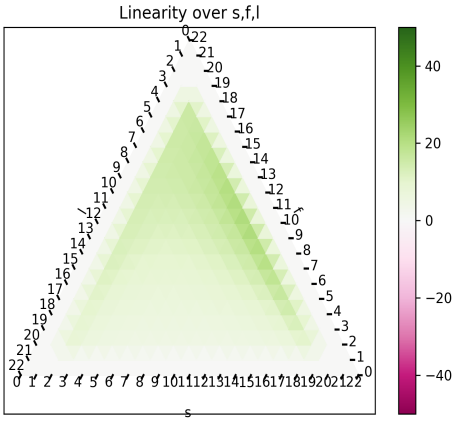

Supplemental Figure S7: This plot shows the difference in initial and end-state linearity with agents with no opponent assessment (acuity variance = ∞). Except in certain edge cases, linearity always increased.

S.7 The number of agents impacts the persistence and extent of linearity

The number of agents also impacts linearity. We explored this in a context where all agents were equal size, thus the "correct" dominance hierarchy would be flat and/or random. We simulated many rounds of contests with few (n=5) and many (n=50) individuals (Figure S6). In both cases, linearity started above random chance and initially increased, however with many individuals, linearity eventually fell back to baseline. This is likely due to the large amount of information individuals received causing them to all (correctly) conclude that they were the same size, where in the low number, low information environment, it was easy to establish relative assessments to enforce linearity.

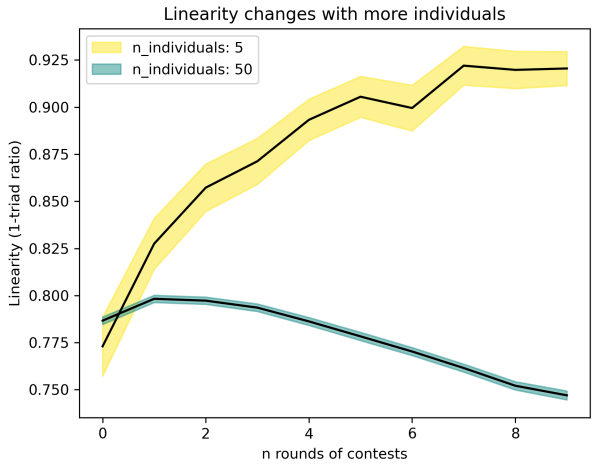

Supplemental Figure S8: Artificial linearity among equally sized individuals is stronger and more persistent in small social groups. This plot shows mean linearity for two simulated dominance networks, one with 5 agents, a second with 50 agents. In both cases, all individuals were the same size ($x_i = 50$). Error bars show SEM over 1000 iterations.

S.8 What drives the experience effect under linear updating

This similarity between Bayesian and linear updating is due to the divergent effects of linear updating on size updating. With linear updating, any individual that is larger—or smaller—than the median size is winning—or losing—the majority of their contests, respectively. This means agent estimates will diverge over time, with all individuals bigger than the median eventually assuming they are extremely large and all below-average individuals concluding they are extremely small. [Make a figure for this]

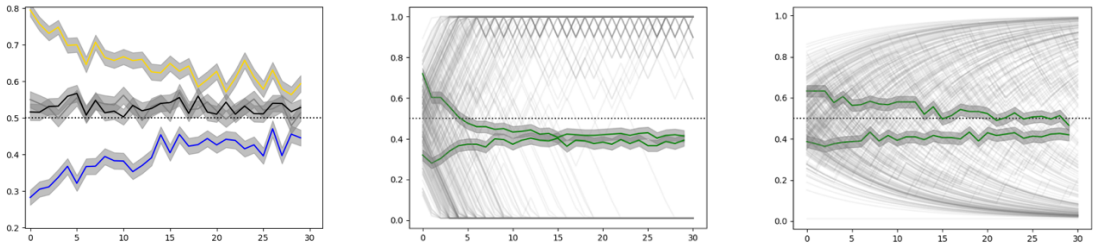

Supplemental Figure S9: Winner-loser affects attenuate with experiences. A. Compared to not updating estimates (black), Bayesian updating shows both a winner (gold) and a loser (blue) effect which attenuates over time. B. Linear updating (green) also shows attenuation over time, even though the actual shift is fixed. This is because individuals' estimates (black lines, plotted as a ratio of maximum possible size) diverge to either the maximum or minimum sizes, since any individual that is not precisely the median size will necessarily either win or lose the majority of contests. C. This same pattern holds even if the shift is a function of the size from min/max size. Unless the shift is extremely large (i.e., 50% of maximum size), a single experience isn't enough to compensate for past experience, and so winner-loser effects are weaker for experienced individuals.

References

Albers, P. C. H., and H. De Vries. 2001. Elo-rating as a tool in the sequential estimation of dominance strengths. Animal Behaviour.

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.

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.

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

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.

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.

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.

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.

Hemelrijk, C. K. 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.

Johnsson, J. I. 1997. Individual recognition affects aggression and dominance relations in rainbow trout, Oncorhynchus mykiss. Ethology 103.
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.
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.
McDonald, A. L., N. W. Heimstra, and D. K. Damkot. 1968. Social modification of agonistic behaviour in fish. Animal Behaviour 16.
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.
Neumann, C., J. Duboscq, C. Dubuc, A. Ginting, A. M. Irwan, M. Agil, A. Widdig, et al. 2011. Assessing dominance hierarchies: Validation and advantages of progressive evaluation with Elo-rating. Animal Behaviour.
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.
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.
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.