Heterogeneous responsiveness to environmental stimuli*

Jerome Cavailles[1], Christoph Kuzmics[2], and Martin Grube[1]

[1]Institute of Biology, University of Graz, Graz, Austria
[2]Institute of Economics, University of Graz, Graz, Austria

January 26, 2023

Keywords: equilibrium, ESS, behavioral ecology, individual differences

Abstract

Individuals of a species cope with environmental variability through behavioral adjustments driven by individuals' responsiveness to environmental stimuli. Three key empirical observations have been made for many animal species: The *coexistence* of different degrees of responsiveness within one species; the *consistency* of an individual's degree of responsiveness across time; and the *correlation* of an individual's degree of responsiveness across contexts. Taking up key elements of existing approaches, we provide one unifying explanation for all three observations, by identifying a unique evolutionarily stable strategy of an appropriately defined game within a stochastic environment that has all three features. Coexistence is explained by a form of negative frequency dependence. Consistency and correlation is explained through potentially small, individual, differences of states animals have and the resulting differential advantages they can get from it. Our results allow us to identify a variety of testable implications.

*We are grateful to Adam Clark, Bernat Corominas-Murtra, Michael Greinecker, Josef Hofbauer, John McNamara and Jorge Peña as well as seminar participants at the Institute for Advanced Studies, Toulouse, the participants at COLIBRI focus workshop on Game Theory in Biology in Graz and the members of the field of excellence initiative COLIBRI at the University of Graz for helpful comments and suggestions.

1 Introduction

The paper in a nutshell A first step to understand how resilient different species are to an increase in the variability of their environment is to understand the driving mechanisms behind their ability to deal with the already volatile environments that they have always lived in. We build a game-theoretic model to identify potential mechanisms leading to an evolutionary stable strategy that exhibits the key empirical observations that have been made in the literature on the responsiveness to environmental stimuli in a stochastic environment. The model also delivers additional testable implications. The model provides specific quantitative frequency predictions of the distribution of environmental responsiveness as a function of a few parameters of the environment. Owing to the stylized nature of the game theoretic model, these quantitative predictions should be taken with a grain of salt, but could at least be the starting point for additional exploration. The model, however, also provides a range of robust qualitative predictions. The model, in general, predicts a continuum of different degrees of environmental responsiveness. It also predicts that the exact stochastic nature does not affect the distribution of environmental responsiveness. This implies that any potential changes to the stochastic process that generates these individuals' environment does at least not affect the distribution of these individuals' degree of responsiveness. Finally, when there is idiosyncratic noise in the individuals' observation of their environment, the model predicts that the equilibrium responsiveness increases when the observation noise increases. More observation noise, in some sense, forces individuals to overreact to environmental stimuli.

Building blocks from the existing literature It is well documented that individuals differ in the degree of responsiveness to external stimuli, a phenomenon sometimes referred to as behavioral plasticity, see e.g., Mitchell and Houslay [2021]. This difference in environmental responsiveness constitutes the main characteristic of personalities, and personalities have been observed in more than 100 species, see e.g., the survey by Nettle and Penke [2010]. Three key observations have been made for many animal species, as, for instance, highlighted by Wolf et al. [2007] (see also Bell et al. [2009]): The *coexistence* of different degrees of responsiveness within one species; the *consistency* of an individual's degree of responsiveness across time; and the consistency, or often referred to as *correlation*, of an individual's degree of

responsiveness across contexts.[1]

A few theoretical approaches explain one or more of these three observations. The theory of biological sensitivity to context, as in Pluess [2015], Boyce [2016], and Ellis and Boyce [2008] explains the coexistence of different degrees of environmental responsiveness with differences in individuals' experiences in their early development, where different experiences lead to different behavior. The theory of differential sensitivity as in Ellis et al. [2011] interprets the difference of behaviour as a way to hedge future generations against the uncertainty in the environment, recently formalised by Bergstrom [2014] and Frankenhuis et al. [2016]. Since the future is inherently unpredictable, parents have offspring with different degrees of environmental responsiveness so that for every environment there are at least some offspring that are able to cope with the environment.[2] A third theoretical approach is built around the idea of "negative frequency dependence:" The more individuals are responsive to environmental stimuli the less the benefits of being responsive. Negative frequency dependence is a cornerstone for explanations of the coexistence of different degrees of environmental responsiveness in the seminal models of Wolf et al. [2008], Wolf et al. [2011], Wolf and McNamara [2012]. See Dingemanse and Wolf [2010] for a review of earlier models. While negative frequency dependence is able to explain coexistence, consistency and correlation are often explained by an individual's state (e.g., morphology, phenotype, size, etc.), as in Wolf and Weissing [2010] and Dingemanse and Wolf [2010].[3] However, the recent meta analysis by Niemelä and Dingemanse [2018] shows a weak link between state and personalities (individuals' state can only explain between 3 and 8% of the personality differences).

We study individual responsiveness to environmental stimuli for the specific problem of foraging from multiple food sources, which allows us to additionally build on the existing game-theoretic literature on the *ideal free distribution* of Fretwell [1969], see e.g., the survey of Křivan et al. [2008]: individuals allocate themselves proportionally to the amount of food available at each food source. We enrich these game-theoretic models by embedding

[1] The "suite of correlated behaviors [...] reflecting the individual consistency across [...] situations" has been referred to as a "behavioral syndrome" by Sih et al. [2004].

[2] The two theories are not exclusive. Pluess [2015], also Mitchell and Houslay [2021] argue in favor of an integration of the two theories.

[3] Wolf et al. [2008] explain consistency and correlation with a "positive-feedback mechanism:" responsiveness is less costly for individuals that have been responsive before. Wolf and McNamara [2012] explain consistency and correlation by small variations of individuals' metabolism (which is a form of state).

them in a stochastic environment, which allows us to incorporate the salient features of the three approaches mentioned above. We do this in a few steps of varying complexity.

Modelling strategy and main results The basic model (in Section 2) is sketched in Figure 1. To illustrate the model we use fish-feeding birds as an example, with the simplifying condition that they do not show any social behavior (such as flocking or swarming). Individuals have to choose to forage from one of two food sources, one providing a fixed amount of food, the other a random amount of food. All individuals who go to the same food source are assumed to share the available food there equally. In the basic model the random food source can only have two possible levels of food availability.

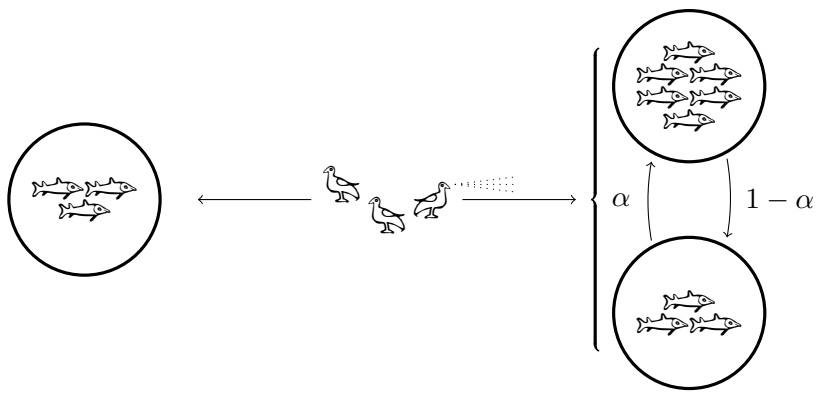

Figure 1: A graphical sketch of the basic model. The birds are the players, who choose which food source to go to: fixed food source A on the left or stochastic food source B on the right, with α the probability of high food availability. The bird's scanning indicates that players make their decision based on observing the food availability at the stochastic source.

Our first main result (Proposition 1) is that, under a mild condition, the game always has a unique symmetric Nash equilibrium strategy, and this strategy exhibits coexistence of responsive and non-responsive individuals.[4]

[4]We use the definition of an ESS as extended from that of Smith and Price [1973] to symmetric n-player games by Palm [1984], see also Broom et al. [1997]. In Proposition 2 we show that our games are *stable games* in the sense of Hofbauer and Sandholm [2009]. This implies that the unique symmetric Nash equilibrium of these games is asymptotically stable under most plausible behavioral adjustment dynamics, such as the replicator dynamics

This coexistence is in terms of the equilibrium being in completely mixed strategies. This has two possible interpretations. Either all individuals truly randomize between being responsive and not being responsive. Under this interpretation we would not obtain consistency and correlation. Or, alternatively, the individuals playing the game are randomly chosen from a larger population to play this game and an appropriate fraction of individuals is responsive and another is non-responsive. This could be interpreted to at least allow for consistency and correlation. In Section 4 we introduce small perturbations to the game theoretic model be allowing an individual's food source preference, and the individual's (small) cost of being responsive, to be a little bit idiosyncratic. This is essentially the idea of Harsanyi [1973] *purification* and the idea of *threshold decisions* of McNamara and Houston [2005]. The mixed strategy equilibrium (or ESS) corresponds in the so modified game to an equilibrium in which every individual actually uses a pure strategy of either being responsive or not. This pure strategy choice depends to some extent on this individual's idiosyncratic preference or cost and would thus be stable across time and, to a lesser extent, also across contexts.

We not only show that there is a unique symmetric Nash equilibrium with the desired features, but also provide an analytic expression for the equilibrium pure strategy frequencies as a function of the parameters of the problem. This allows us to derive additional testable predictions of our model. Acknowledging that our model is a highly simplified account of reality, some of these predictions may well hold beyond the narrower confines of our model. In particular, and perhaps most striking: the unique equilibrium does not change with the stochasticity of the environment, at least when the cost of cognition is negligible. In other words, the equilibrium does not depend on the probability of the state of food availability at the random food source. This implies that changes to the stochastic process of the environment may not push behavior out of equilibrium. Individuals' strategies are already sufficiently complex to allow essentially immediate and automatic adaptation to such changes. Of course, this does not imply that individuals are not affected if, for instance, high food availability becomes rarer. Only their strategy is unaffected, not necessarily the amount of food they can consume.

Another testable prediction can be derived from an extension of our

of Taylor and Jonker [1978], the Smith dynamics [Smith, 1984], the projection dynamics [Nagurney and Zhang, 1997], as well as all dynamics of the class of target dynamics (among them the BNN [Brown and Von Neumann, 1950], Best response [Gilboa and Matsui, 1991] and Logit [Fudenberg and Levine, 1998] dynamics).

basic model, presented in Section 5, in which individuals only receive (private) noisy information about the state of food availability. At least for small amounts of noise, the equilibrium responsiveness increases when the noise increases. More noise, in some sense, forces individuals to overreact to environmental stimuli. A final testable prediction, again derived from an extension of our basic model and presented in Section 6, is that for more general distributions of food availability, one would expect to see a continuum of degrees of responsiveness to environmental stimuli. See e.g., Lionetti et al. [2018] for empirical support for this finding.

2 The basic model

In this section we present the simplest possible model of interest for our problem. Each of $n < \infty$ individuals can go to one of two food sources A or B. Food source A has a fixed amount of food normalized to n units of nutrition to facilitate an easier comparison when we vary the number of individuals n. Food source B has a stochastic amount of food $n \cdot X$, with X drawn from a Bernoulli distribution with $X = \eta$ with probability α and $X = \lambda$ with probability $1 - \alpha$.

Before making their choice of food source, individuals can, in principle, inform themselves about the state of food source B; individuals can choose to learn whether $X = \lambda$ or $X = \eta$. Individuals, thus, have six strategies at their disposal in this simple model. They can choose to be informed or to be uninformed. If uninformed they can then choose to go to food source A or B. If informed they have the choice to ignore their information and go to A or B regardless of the information they have received, or they can react to the information in one of two ways. They can be, what we term, *responsive* by going to A when $X = \lambda$ and going to B when $X = \eta$, or, what we term *counter-responsive* by doing the opposite.

We assume that the choice of becoming informed bears an arbitrarily small cost $c > 0$, where arbitrarily small means that we ultimately investigate the limit case when c tends to zero. The assumption of positive costs implies that the two strategies of getting informed and then ignoring the information are strictly dominated by the strategy of not getting informed and going to the same food source. By virtue of saving on the small positive cost c the latter strategies provide a strictly higher payoff in both states $X = \lambda$ and $X = \eta$ regardless of what the other individuals do. We are thus left with four pure strategies. We denote the set of strategies by $S = \{A, B, R, C\}$, for always going to food source A, always to B, being responsive, and being

counter-responsive, respectively.

We allow individuals to choose randomly. We come back to the interpretation of random choice in Section 4. Let $\Delta(S)$ denote the set of all mixed strategies, that is the set of all probability distributions over S. A pure strategy $s \in S$ can be identified as the mixed strategy that attaches probability 1 on pure strategy s.

We assume throughout the paper that all individuals who go to the same food source share the available food at this source equitably. This implies that the payoff to an individual who goes to one of the two food sources only depends on the number of other individuals, k, that go to food source A (which implies that $n-1-k$ others go to food source B). The payoff to an individual who goes to food source A is given by $\frac{n}{k+1}$; the payoff to an individual who goes to B is given by $\frac{nX}{n-k}$.

Given that individuals can choose mixed strategies, that is can choose randomly, we need to compute individuals' expected payoffs given these random choices. To do so, consider an arbitrary individual who is facing that all other $n-1$ individuals choose a given (mixed) strategy $\sigma \in \Delta(S)$ with $\sigma(s)$ the probability that pure strategy s is chosen. Denote by N_s the random variable that is the number of opponents who end up choosing pure strategy $s \in S$ (given the probability of choosing s is $\sigma(s)$). The tuple (N_A, N_B, N_R, N_C) is then multinomially distributed with parameters $n-1$ and probability vector $(\sigma(A), \sigma(B), \sigma(R), \sigma(C))$.

Given σ, let $R_{A\lambda}$ denote the random variable that is the food share available at food source A in state λ. Let food shares $R_{A\eta}$, $R_{B\lambda}$, $R_{B\eta}$, be defined analogously. These (random) food shares are given by

$$\begin{array}{rcl} R_{A\lambda} & = & \frac{n}{N_A+N_R+1} \\ R_{A\eta} & = & \frac{n}{N_A+N_C+1} \\ R_{B\lambda} & = & \frac{\lambda \cdot n}{N_B+N_C+1} \\ R_{B\eta} & = & \frac{\eta \cdot n}{N_B+N_R+1}. \end{array}$$

We can, then, express an individual's expected payoffs from choosing pure strategy $s \in S$, when all others use mixed strategy σ, as follows.

$$\begin{array}{rclcl} u(A,\sigma) & = & \alpha \mathbf{E}\left[R_{A\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{A\lambda}\right] \\ u(B,\sigma) & = & \alpha \mathbf{E}\left[R_{B\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{B\lambda}\right] \\ u(R,\sigma) & = & \alpha \mathbf{E}\left[R_{B\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{A\lambda}\right] \quad - \quad c \\ u(C,\sigma) & = & \alpha \mathbf{E}\left[R_{A\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{B\lambda}\right] \quad - \quad c. \end{array} \tag{1}$$

In the above expression, $\mathbf{E}$ denotes the expectation with respect to the

randomness created by mixed strategy σ.[5]

We extend an individual's payoff function to mixed strategies by taking expectations. That is

$$u(\sigma', \sigma) = \sum_{s \in S} u(s, \sigma)\sigma'(s).$$

A strategy $\sigma \in \Delta(S)$ is a *symmetric Nash equilibrium*, see Nash [1950], if $u(\sigma, \sigma) \geq u(\sigma', \sigma)$ for all $\sigma' \in \Delta(S)$.

3 Results for the basic model

The following proposition (with proof in the appendix) is the main result for the basic model for the limit case where the cost of cognition (or being responsive) is negligible. It proves that there is a unique symmetric Nash equilibrium of this game, shows that in this equilibrium there is heterogeneity in individual's responsiveness to environmental stimuli, and provides an exact analytic expression for the frequency of responsive individuals. This equilibrium responsiveness frequency, finally, is shown to be independent of the stochasticity α in the environment.

Proposition 1. *Under the assumption that $\frac{1}{n} < \lambda < \eta < n$, the game in the basic model, in the limit as cognition cost $c \to 0$ and the number of individuals $n \to \infty$, has a unique limiting symmetric Nash equilibrium, in which $\sigma(A) = \frac{1}{1+\eta}$, $\sigma(B) = \frac{\lambda}{1+\lambda}$, $\sigma(R) = \frac{\eta-\lambda}{(1+\lambda)(1+\eta)}$, and $\sigma(C) = 0$.*

The intuition behind the key insight of this result is as follows. First, we notice that no pure strategy can be an equilibrium. If everyone always goes to food source A, it would be best to go to food source B. If everyone goes to B, it would be best to go to A. If everyone is responsive, then it would be best to be counter-responsive, so as to always be at the food source where no one else is and not have to share food at all. If everyone is counter-responsive, for the same reason it is best to be responsive. One can then show, in fact, that in a Nash equilibrium there must be a mix of strategies A, B, and R (and, interestingly, no C). This means that, in equilibrium, we have the coexistence of responsive and unresponsive individuals. This also implies that all individuals must be indifferent between always going to A, always going to B, and to being responsive. This can only be achieved

[5] For instance, $\mathbf{E}\left[R_{A\lambda}\right] = \sum_{k_A=0}^{n-1} \sum_{k_R=0}^{n-1-k_A} \frac{n}{k_A+k_R+1} \frac{(n-1)!}{k_A!k_R!(n-k_A-k_R)!} \sigma(A)^{k_A} \sigma(R)^{k_R} (1-\sigma(A)-\sigma(R))^{n-k_A-k_R}$.

if the realized food share at each food source is the same, regardless of the state of the food availability at the random food source. The exact equilibrium frequencies can then be derived from this argument. Finally, this finding then also explains why the frequency of responsive individuals does not depend on the exact stochastic nature of the environment.

It is well known that a necessary condition for a strategy to be an evolutionarily stable strategy (ESS) in the sense of Smith and Price [1973], and for our context in the sense of Palm [1984], is that the strategy in question is a symmetric Nash equilibrium strategy. It is also well known, see e.g., Nachbar [1990], that Nash equilibria are the only candidates for asymptotically stable rest points under most deterministic behavioral adjustment (or evolutionary) dynamics, with the replicator dynamics of Taylor and Jonker [1978] the first and most prominent example. The following result shows that the unique symmetric Nash equilibrium that we identified in Proposition 1 is indeed an ESS as well as asymptotically stable under many behavioral adjustment dynamics.

Proposition 2. *Under the assumption that $\frac{1}{n} < \lambda < \eta < n$, the unique equilibrium of the game in the basic model, in the limit as cognition cost $c \to 0$ and the number of individuals $n \to \infty$, is an ESS, is asymptotically stable under the replicator dynamics [Taylor and Jonker, 1978], the Smith [Smith, 1984] and the projection dynamics [Nagurney and Zhang, 1997], as well as the class of target dynamics which comprises, among others, the BNN [Brown and Von Neumann, 1950], Best-response [Gilboa and Matsui, 1991] and Logit dynamics [Fudenberg and Levine, 1998].*

The key insight to prove this Proposition is that all games in our class of games are stable games in the sense of Hofbauer and Sandholm [2009], as we prove in the Appendix. All results then follow from this fact.

Some of these findings can be seen in Figure 2, which depicts the phase diagram of the replicator dynamics for a variety of different parameter specifications. The only difference between Figures 2.a and 2.b is that the stochasticity parameter α changes. This has no effect on the equilibrium itself, but does affect somewhat the out of equilibrium dynamics. It can also be seen that not only is the unique equilibrium asymptotically stable, but also in fact a global attractor under the replicator dynamics: All solution paths eventually converge to the equilibrium.

We finally highlight that, if the cost of cognition, c, is non-negligible, then the equilibrium frequencies not only depend on this cost, but also depend on α, the stochastic nature of the environment. By the known fact,

see e.g., Kohlberg and Mertens [1986], that the (symmetric) Nash equilibrium correspondence is upper-hemi-continuous in the space of games, small changes to c and α can, however, only lead to small changes in the equilibrium frequencies.

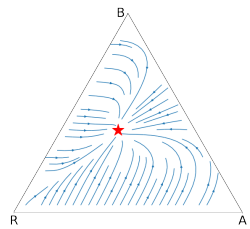

(a) $\alpha = 0.7, \lambda = 0.7, \eta = 3$

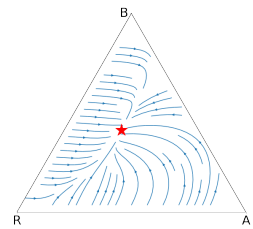

(b) $\alpha = 0.4, \lambda = 0.7, \eta = 3$

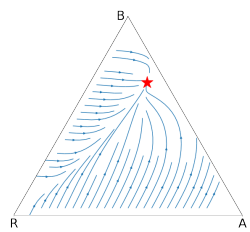

(c) $\alpha = 0.4, \lambda = 2, \eta = 3$

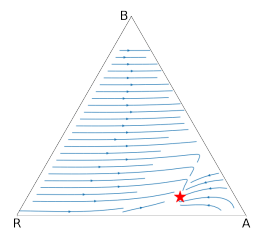

(d) $\alpha = 0.1, \lambda = 0.1, \eta = 0.5$

Figure 2: Phase diagram of the replicator dynamics for different parameter values. The cost of cognition c is always zero. The equilibrium is denoted by a red star.

4 Consistency and Correlation

Suppose that the n individuals play the same game given in our basic model over and over again for many periods of time. Suppose that, at every point in time, they play the unique symmetric equilibrium given in Proposition 1. An outside observer would note that when the amount of food available at B is high ($X = \eta$) more individuals are to be found at source B than when it is low ($X = \lambda$). They would also observe that the food share each individual receives is the same regardless of which source the individuals go to. The outside observer would conclude that some individuals must

be responsive to the stochastic food availability at source B. But if they were to trace each individual, they would realize that each individual is sometimes responsive and sometimes not responsive. There is no internal consistency in the individuals' behavior. Each animal randomizes at each point in time. This is no surprise, because the unique equilibrium is in mixed (or randomized) strategies. However, we must point out that this is inconsistent with empirical findings.

A slight change to the basic model is able to accommodate the common empirical finding that over time typically the same individuals are responsive. This modification is based on the idea of *purification* of Harsanyi [1973], which is very similar to the idea of *threshold decisions* as provided in McNamara and Houston [2005]. The idea is that individuals differ a little bit in terms of their personal preferences and actually make a pure strategy choice that is, however, dependent on their own personal preferences that only they themselves know. As a consequence, while the equilibrium looks mixed to other individuals, each individual actually plays a pure strategy. We adapt the model according to this idea by replacing the payoff function u of the original model with a slightly perturbed payoff function v_θ that is essentially equal to u plus a small idiosyncratic (individual-specific) preference or perturbation term:

$$\begin{array}{rcll} v_\theta(A,\sigma) & = & u(A,\sigma) & + \quad \theta_A \\ v_\theta(B,\sigma) & = & u(B,\sigma) & + \quad \theta_B \\ v_\theta(R,\sigma) & = & u(R,\sigma) & + \quad \alpha\theta_B + (1-\alpha)\theta_A + \theta_R \\ v_\theta(C,\sigma) & = & u(C,\sigma) & + \quad \alpha\theta_A + (1-\alpha)\theta_B + \theta_R, \end{array}$$

where the vector $\theta = (\theta_A, \theta_B, \theta_R)$ is i.i.d. drawn from some arbitrary full support continuous joint distribution F (with density f) over $\Theta = [-\epsilon, \epsilon]^3$, for a small $\epsilon > 0$. It is assumed that an individual's realized θ is that individual's private information, unknown to other individuals. We have deliberately chosen the same preference perturbation θ_R for pure strategies C and R, as it seems more natural to have an idiosyncratic perturbation of the cost of being responsive rather than for how one is responsive. However, it does not matter what we assume for pure strategy C as long as the payoff perturbation is small, as pure strategy C provides a strictly lower payoff than the other three strategies in the equilibrium given in Proposition 1, and small payoff perturbations cannot change that.

Harsanyi [1973] has shown that almost any equilibrium of a complete information game, such as our basic game, is such that for any nearby incomplete information game with payoff perturbations given by the joint

distribution F has a nearby equilibrium and that this nearby equilibrium is essentially in pure strategies.[6] In such a nearby equilibrium there is a parameter region for $\theta \in \Theta$ for which an individual strictly prefers to play A, another region for which an individual strictly prefers to play B, and a final region in which an individual strictly prefers to play R. The set of θ's for which an individual is indifferent between two or three of the three strategies has measure zero. See Appendix B for more details. Finally, such purified equilibria can also shown to be dynamically stable under a suitably defined behavioral adjustment dynamics as in Ely and Sandholm [2005], see Sandholm [2007].

Suppose now that the n individuals play the resulting equilibrium of the same slightly perturbed game repeatedly over many time periods. It is then a question of whether the perturbation parameters θ remain the same for each individual over time or not. Suppose that they do. Then an outside observer would not only observe all that the observer would have observed that we described above, but also that it would be the same individuals who always go to food source A, the same individuals who always go to food source B, and the same who are responsive.

This model is also flexible enough to generate a strong consistency over time and a weaker, but some, consistency across contexts, depending on how these consistencies are interpreted. Consider the bird example again. One could imagine that θ_R is an individual bird's specific parameter that does not change over time nor across contexts. On the other hand the parameters θ_A and θ_B might be constant for one season, but could be different in another season, when the bird's nest location (or the location of the food sources) changes.

As an example of why the perturbed model may be appropriate for our purposes, consider birds who every day have to decide to go to food source A or B from their nesting place. Then the location of their nesting place gives rise to their θ. An approach could be that θ_A and θ_B are proportional to the distance that the bird's nest is from the two food sources, respectively, while θ_R could be more of a personal characteristic of the bird, measuring how much/less cognitively able this bird is relative to other birds.

One could obtain equilibrium purification even by introducing a payoff-irrelevant personal and privately know characteristic, such as an individual's prior experiences in life, with individuals playing different pure strategies

[6] Harsanyi [1973] calls such equilibria *regular* and shows that all finite complete information games have a regular equilibrium. As our equilibrium is unique it must be regular.

depending on their personal prior experiences. This means that, as pointed out e.g., in Wolf et al. [2007], Wolf and Weissing [2010] and Dingemanse and Wolf [2010], the purification threshold could also be based on an individual's state or their life history.

5 Imperfect private signals of food availability

In our basic model, individuals can learn the state of food availability at the food sources perfectly. In this section we study how the results change if this learning is imperfect. To do so we suppose that each individual, when they learn, receive a noisy signal about the actual level of food availability at food source B. Individuals i receive conditionally independent (and identically distributed) signals $s_i \in \{l, h\}$ such that $P(s_i = h|X = \eta) = P(s_i = l|X = \lambda) = 1 - \epsilon$, with $\epsilon < \frac{1}{2}$. In words, in the high state η the high signal h is more likely than the low signal l and in the low state λ the low signal l is more likely than the high signal h. The signal is, thus, informative about the true state.

In this case, (random) food shares are given by

$$\begin{array}{rcl} R_{A\lambda} & = & \frac{n}{N_{A\lambda}+1} \\ R_{A\eta} & = & \frac{n}{N_{A\eta}+1} \\ R_{B\lambda} & = & \frac{\lambda \cdot n}{N_{B\lambda}+1} \\ R_{B\eta} & = & \frac{\eta \cdot n}{N_{B\eta}+1}, \end{array}$$

where, given strategy σ (used by all opponents), $N_{A\lambda}$ follows a binomial distribution with $n-1$ trials and success probability $\sigma(A) + (1-\epsilon)\sigma(R) + \epsilon\sigma(C)$. Similarly, $N_{A\eta}, N_{B\lambda}, N_{B\eta}$ are also binomially distributed with $n-1$ trials and success probabilities $\sigma(A) + (1-\epsilon)\sigma(C) + \epsilon\sigma(R)$, $\sigma(B) + (1-\epsilon)\sigma(R) + \epsilon\sigma(C)$, and $\sigma(B) + (1-\epsilon)\sigma(C) + \epsilon\sigma(R)$, respectively.

Payoffs are as described in Equation 1. All the arguments of the proof of Proposition 1 go through and, ultimately, we obtain that the game has a unique symmetric Nash equilibrium, σ, that is also an ESS (and all other results of Proposition 2 apply) with

$$\left\{ \begin{array}{rcl} \sigma(R) & = & \frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)} \\ \sigma(A) & = & \frac{1}{1+\eta} - \epsilon\frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)} \\ \sigma(B) & = & \frac{\lambda}{1+\lambda} - \epsilon\frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)}. \end{array} \right.$$

The key new insight is that the higher the error probability ϵ the higher

the fraction (or probability) of responsive individuals in equilibrium.

6 General distributions of food source availability

Recall that, in the basic model food source B is assumed to be Bernoulli distributed (i.e., with only two possible levels of available food at that source). In this section we consider an arbitrary distribution for the food availability at food source B.

Let X, the available quantity of food at food source B, be distributed according to some distribution with cdf F with everywhere positive density f on the interval $\chi = [x_L, x_H]$ with $0 \leq x_L < x_H \leq \infty$. To make the analysis tractable we simplify the model in two ways. First, we set the cost of being responsive, $c = 0$.[7] Second, we assume that all individuals learn the value of X, and allow individuals to only use monotone strategies: An individual's strategy can be described by a cutoff value $y \in \chi$ such that the individual goes to food source A if and only if $x < y$. Otherwise the individual goes to food source B. This implies that the strategy space is identical to χ and the set of mixed strategies is the set $\Delta(\chi)$ of all probability distributions over χ. A fully mixed symmetric Nash equilibrium strategy, which can be described by a cdf G on χ must satisfy that any individual is indifferent between using any pure strategy in χ, i.e., between using any cut-off $y \in \chi$.

We then get the following result.

Proposition 3. *In the model of this section, for any n, there is a unique completely mixed symmetric equilibrium. In the limit as n tends to infinity, the equilibrium probability that an individual uses cut-off reactivity y is given by the cdf $G(y) = \frac{y}{1+y}$, with $G(x_L) = \frac{x_L}{1+x_L}$ the probability of an individual always going to food source B, $1 - G(x_H) = \frac{1}{1+x_H}$ the probability of an individual always going to food source A, and $G(y) - G(x) = \frac{y-x}{(1+x)(1+y)}$ the probability that an individual adopts a degree of responsiveness in the interval $[x, y]$.*[8]

This finding is consistent with those in the basic model. For example, the strategy called B in the previous model is here similar of choosing the

[7] In a setting in which individuals can have different degrees of responsiveness, as we have here, one might want to assume that the cost of responsiveness varies with the degree of responsiveness. We shall not pursue this here, however.

[8] The distribution with cdf $G(y) = \frac{y}{1+y}$ is the distribution of a random variable Y such that its reciprocal (or inverse) $\frac{1}{Y}$ has exactly the same distribution. One could call G the inversion invariant distribution.

cut off x_L (since x_L is the minimum possible value of the stochastic source). The equilibrium frequency of this strategy in the basic model is given by $\frac{\lambda}{1+\lambda}$, which is equivalent to $\frac{x_L}{1+x_L}$, as x_L is the smallest possible value of the stochastic source.

7 Discussion

We built a stylized game theoretic model of foraging behavior in a stochastic environment. For every parameter specification within certain bounds, this model has a unique symmetric Nash equilibrium, that is also the unique ESS and asymptotically stable under a variety of evolutionary dynamics. This equilibrium has the three key features identified in the literature of *coexistence* of differing degrees of environmental responsiveness, *consistency* of individual environmental responsiveness over time, and *correlation* of individual environmental responsiveness across contexts.

By explicitly studying the phenomenon of heterogeneous responsiveness to environmental stimuli in a foraging setting, we are able to identify the push towards the ideal free distribution of Fretwell [1969], satisfied in the equilibrium of our game, as a possible driving force of this heterogeneity. We derive explicit analytical expressions for the equilibrium frequencies of responsive and non-responsive behavior, at least when the cost of cognition (needed to respond to environmental stimuli) is negligible or at least relatively small. This allows us to perform comparative statics as to how the equilibrium frequencies change when some of the model parameters change. For instance, we find that, at least when cognition costs are negligible, the exact stochastic nature of the environment does not affect the equilibrium. This finding suggests that changes in the stochastic environmental process would at least not be so disruptive as to push behavior out of equilibrium. Put differently, equilibrium strategies are already complex enough to allow for automatic adaptation to such changes in the stochastic environmental process. In the remainder of this section we discuss some of the limitations of our approach.

Cost of cognition We have only explored the case of zero, and by a continuity argument, also of small cost of cognition. For larger cost of cognition, generally, equilibrium behavior will depend on the stochastic nature of the environment and the equilibrium will not satisfy the ideal free distribution. While it would be interesting to do so, we have not explored this issue. We feel, however, that for most species, it is not unreasonable to assume a

relatively small cost of cognition.

Noisy information Another insight that we can derive from an extension of our model is that the higher the noise in the environmental stimuli the more responsive individuals become in equilibrium. This is under the assumption of individuals receiving private and stochastically independent noisy information about the state of the environment. We have not explored the perhaps sometimes more interesting case of correlated information, such as all individuals receiving the same public information. In such a setting, the ideal free distribution would at best hold in expectation, and there would be a positive variance of food *share* availability at the random source. Another, empirically relevant, informational setting is one where not all individuals receive the same quality of information (perhaps not all are equally close to the source of information). One would then expect individuals to infer additional information about the state of the environment from other individuals' behavior. If, for instance, there are many birds flying out to a specific point at see, another bird might follow based on the idea that there is information in that behavior. This will certainly be the case for socializing birds, which display behavior of forming flocks and swarms. Such behavior would add another layer of complexity to the game.

The number of food sources Our model only has two food sources. This keeps the analysis mathematically tractable, but comes at the cost of a possible oversimplification. Introducing an arbitrary number of stochastic food sources would, indeed, make the analysis more complicated. However, given our results, one would conjecture that in any (evolutionary stable) equilibrium of such a modified game, the ideal free distribution holds, at least when costs of cognition are negligible: all food sources would have equal food shares, and this would be true for all states. This alone would already imply that equilibrium behavior would not depend on the exact stochastic nature of the environment. This would also imply that any (evolutionary stable) equilibria would again satisfy coexistence, and for slightly perturbed models, consistency and correlation. The only difference is that one would not necessarily expect a unique equilibrium and it would be harder to characterize these explicitly.

Generalizing We focused here on foraging choice as a concrete setting in which one would expect the coexistence, consistence and correlation of different responsiveness to external stimuli. However, our results could pos-

sibly be generalized to different context where there is a resource to share among individuals, with a resource distributing at different point, and some of them stochastic. Such contexts are social interactions, mating behaviour, division of labour [Dall et al., 2012], space-use [Spiegel et al., 2017], or niche specialisation [Dall et al., 2012, Montiglio et al., 2013, Bolnick et al., 2003]. Those last studies shows similar concepts (state dependence, frequency dependence, social awareness, environmental heterogeneity) applied to niche specialisation. In particular, increasing evidence show a link between specialisation and personality. It is hypothesized that personality implies specialization [Toscano et al., 2016, Harris et al., 2020] or the other way around [Bergmüller and Taborsky, 2010, Montiglio et al., 2013]. Finally and more speculatively, our results might even be translatable to the issue of stem cell differentiation, which might arise from competition over resources, see Jörg et al. [2019].

A Proof of Proposition 1

We first show that pure strategies A, B, and R, cannot be equilibria for any $c > 0$ and for any n given the assumption that $\frac{1}{n} < \lambda < \eta < n$. Suppose everyone plays A. Then everyone gets a payoff of $\frac{1}{n}$. An individual deviating to A would at worst get λ and since $\lambda > \frac{1}{n}$ this is a profitable deviation. Suppose everyone plays B. Then everyone gets the same payoff of at most $\frac{\eta}{n} < 1$ (if $X = \eta$ and less if $X = \lambda$). Deviating to B would yield a payoff of 1. Suppose, finally, that everyone plays R. Then everyone gets a payoff of either $\frac{1}{n}$ (if $X = \lambda$) or $\frac{\eta}{n}$ (if $X = \eta$), minus c, and then a deviation to B gives a higher payoff of λ when $X = \lambda$ and the same payoff if $X = \eta$, and thus, a strictly higher expected payoff.

Next we show that there cannot be a symmetric Nash equilibrium that attaches non-zero probability to both pure strategies R and C. If this were the case, then necessarily $u(C, \sigma) = u(R, \sigma)$ and, by the fact that $u(A, \sigma) + u(B, \sigma) = u(R, \sigma) + u(C, \sigma) - 2c$ for all $\sigma \in \Delta(S)$, either pure strategy A or B (or both) provides a strictly higher payoff than both R and C and would thus offer a beneficial unilateral deviation.

There cannot be a symmetric Nash equilibrium in which strategy C is used with non-zero probability. To see this, we show that whenever C is used R must be used with non-zero probability. As we just showed that there cannot be an equilibrium with both C and R, this then proves that there is no equilibrium in which C is used.

We prove this by contradiction. Suppose, that $\sigma(R) = 0$ and $\sigma(C) > 0$

for an equilibrium $\sigma \in \Delta(S)$. We will show that, in this case, either strategy A or B will provide a strictly better payoff than C does. Noting that $N_R = 0$, we can write the expected payoffs to the three strategies in question as

$$\begin{array}{rcccccc} u(A,\sigma) & = & \alpha\mathbf{E}\left[\frac{n}{N_A+N_C+1}\right] & + & (1-\alpha)\mathbf{E}\left[\frac{n}{N_A+1}\right] & & \\ u(B,\sigma) & = & \alpha\mathbf{E}\left[\frac{\eta\cdot n}{N_B+1}\right] & + & (1-\alpha)\mathbf{E}\left[\frac{\lambda\cdot n}{N_B+N_C+1}\right] & & \\ u(C,\sigma) & = & \alpha\mathbf{E}\left[\frac{n}{N_A+N_C+1}\right] & + & (1-\alpha)\mathbf{E}\left[\frac{\lambda\cdot n}{N_B+N_C+1}\right] & - & c. \end{array}$$

As σ is an equilibrium we must have $u(A,\sigma), u(B,\sigma) \leq u(C,\sigma)$. This implies that

$$\mathbf{E}\left[\frac{\lambda}{N_B+N_C+1}\right] \geq \mathbf{E}\left[\frac{1}{N_A+1}\right],$$

and

$$\mathbf{E}\left[\frac{1}{N_A+N_C+1}\right] \geq \mathbf{E}\left[\frac{\eta}{N_B+1}\right],$$

respectively.

But, since we have $\mathbf{E}N_C > 0$, we have

$$\mathbf{E}\left[\frac{1}{N_A+1}\right] > \mathbf{E}\left[\frac{1}{N_A+N_C+1}\right],$$

and

$$\mathbf{E}\left[\frac{\lambda}{N_B+1}\right] > \mathbf{E}\left[\frac{\lambda}{N_B+N_C+1}\right],$$

which, in combination with the previous two inequalities yields

$$\mathbf{E}\left[\frac{\lambda}{N_B+1}\right] > \mathbf{E}\left[\frac{\lambda}{N_B+N_C+1}\right] \geq \mathbf{E}\left[\frac{1}{N_A+1}\right] > \mathbf{E}\left[\frac{1}{N_A+N_C+1}\right] \geq \mathbf{E}\left[\frac{\eta}{N_B+1}\right],$$

and, thus,

$$\mathbf{E}\left[\frac{\lambda}{N_B+1}\right] > \mathbf{E}\left[\frac{\eta}{N_B+1}\right],$$

which, finally, contradicts our assumption that $\eta > \lambda$.

Next we show that, given the assumption that $\frac{1}{n} < \lambda < \eta < n$, there cannot be any mixed strategy equilibrium with only two pure strategies in its support. We have already shown that there cannot be an equilibrium with C in its support. Suppose there is an equilibrium σ that has a support of $\{A, R\}$. Then $u(A,\sigma) = u(R,\sigma)$. However, $u(B,\sigma) > u(R,\sigma)$, as B provides the same payoff as R when $X = \eta$ and obtains λ when $X = \lambda$, while R obtains $\frac{1}{n} < \lambda$. Suppose there is an equilibrium σ that has a

support of $\{B, R\}$. Then $u(B, \sigma) = u(R, \sigma)$. However, $u(A, \sigma) > u(R, \sigma)$, as A provides the same payoff as R when $X = \lambda$ and obtains 1 when $X = \eta$, while R obtains $\frac{\eta}{n} < 1$.

Suppose there is an equilibrium σ that has a support of $\{A, B\}$. Then $u(A, \sigma) = u(B, \sigma)$. As $\eta > \lambda$ then for sufficiently small costs c, R provides a higher payoff than both A and B. In expectation both food sources A and B have the same food share, but R is able to get a higher payoff by going to B if and only if $X = \eta$.

This, finally, leaves as the only possible equilibria those with support given by $\{A, B, R\}$. Let $r_{yx} = \mathbf{E}[R_{yx}]$ denote the expected food share at source $y \in \{A, B\}$ in state $x \in \{\lambda, \eta\}$. Consider $r_{A\lambda}$. All other expected food shares are calculated analogously. We can write $r_{A\lambda} = n\mathbf{E}[\frac{1}{Y+1}]$, where $Y = N_A + N_R$, and Y, thus, binomially distributed with $n - 1$ tries and success probability $p = \sigma(A) + \sigma(R)$. The first negative moment of such a binomial distribution, $\mathbf{E}[\frac{1}{Y+1}]$, is provided in [Chao and Strawderman, 1972, Equation 3.4] and given by $\frac{1-(1-p)^n}{np}$.

We, thus, obtain

$$\begin{aligned}
r_{A\lambda} &= \tfrac{1-(1-\sigma(A)-\sigma(R))^n}{\sigma(A)+\sigma(R)}, \\
&\text{and analogously} \\
r_{A\eta} &= \tfrac{1-(1-\sigma(A)-\sigma(C))^n}{\sigma(A)+\sigma(C)} \\
r_{B\lambda} &= \lambda \cdot \tfrac{1-(1-\sigma(B)-\sigma(C))^n}{\sigma(B)+\sigma(C)} \\
r_{B\eta} &= \eta \cdot \tfrac{1-(1-\sigma(B)-\sigma(R))^n}{\sigma(B)+\sigma(R)}.
\end{aligned} \tag{2}$$

We need to show that there is a unique equilibrium with support $\{A, B, R\}$. To be indifferent between A and R requires that

$$r_{A\eta} = r_{B\eta} - \frac{c}{\alpha},$$

to be indifferent between B and R requires

$$r_{B\lambda} = r_{A\lambda} - \frac{c}{1-\alpha}.$$

Plugging in the above expressions, and noting that $\sigma(C) = 0$, these two equations lead to

$$\frac{1-(1-\sigma(A))^n}{\sigma(A)} = \eta \frac{1-\sigma(A)^n}{1-\sigma(A)} - \frac{c}{\alpha}, \tag{3}$$

and

$$\lambda \frac{1-(1-\sigma(B))^n}{\sigma(B)} = \frac{1-\sigma(B)^n}{1-\sigma(B)} - \frac{c}{1-\alpha}, \tag{4}$$

respectively. We are interested in the limit as $n \to \infty$ and $c \to 0$ and obtain

$$\frac{1}{\sigma(A)} = \eta \frac{1}{1-\sigma(A)},$$

and

$$\lambda \frac{1}{\sigma(B)} = \frac{1}{1-\sigma(B)},$$

respectively. Solving these two equations yields unique solutions $\sigma(A) = \frac{1}{1+\eta}$ and $\sigma(B) = \frac{\lambda}{1+\lambda}$. This is the only candidate for an equilibrium of this form. As $\sigma(R) = 1 - \sigma(A) - \sigma(B)$ we obtain $\sigma(R) = \frac{\eta-\lambda}{(1+\eta)(1+\lambda)}$. Note, finally, that $\sigma(A), \sigma(B), \sigma(R) \geq 0$, and, thus σ is indeed the unique (mixed strategy) equilibrium.

B Strategies and equilibria in perturbed games

In the perturbed game a strategy is a function $\rho : \Theta \to \Delta(S)$, where ρ attaches to each possible preference type θ a probability distribution over S. Every such strategy ρ induces a strategy $\sigma^\rho \in \Delta(S)$ by $\sigma^\rho(s) = \int_\Theta \rho(\theta)(s) f(\theta) d\theta$ for all $s \in S$, where $\rho(\theta)(s)$ is the probability that $\rho(\theta)$ attaches to s.

An equilibrium strategy ρ is then such that $\rho(\theta)(A) = 1$ if

$$\begin{array}{rcl} u(A,\sigma^\rho) - u(B,\sigma^\rho) & > & \theta_B - \theta_A, \text{ and} \\ u(A,\sigma^\rho) - u(R,\sigma^\rho) & > & \alpha(\theta_B - \theta_A) + \theta_R; \end{array}$$

$\rho(\theta)(B) = 1$ if

$$\begin{array}{rcl} u(A,\sigma^\rho) - u(B,\sigma^\rho) & < & \theta_B - \theta_A, \text{ and} \\ u(B,\sigma^\rho) - u(R,\sigma^\rho) & > & (1-\alpha)(\theta_A - \theta_B) + \theta_R; \end{array}$$

and $\rho(\theta)(R) = 1$ if

$$\begin{array}{rcl} u(A,\sigma^\rho) - u(R,\sigma^\rho) & < & \alpha(\theta_A - \theta_B) + \theta_R, \text{ and} \\ u(B,\sigma^\rho) - u(R,\sigma^\rho) & < & (1-\alpha)(\theta_A - \theta_B) + \theta_R. \end{array}$$

C Evolutionary Stability

We here show that the unique symmetric equilibrium of the basic game is evolutionarily stable as defined, for symmetric n player games, by Palm [1984], see also Broom et al. [1997]. By [Palm, 1984, Proposition 4] a sufficient condition for a strategy to be an ESS in symmetric n player games is as follows.

Lemma 1 (Palm [1984])**.** *A (mixed) strategy σ of a symmetric n-player normal form game is an ESS if*

1. $u(\sigma,\sigma) \geq u(\sigma',\sigma)$ *for all* $\sigma' \in \Delta(S)$ *and*
2. *if* $u(\sigma',\sigma) = u(\sigma,\sigma)$ *for some* $\sigma' \neq \sigma$ *then* $u(\sigma',\sigma') < u(\sigma,\sigma')$.

Proposition 4. *The equilibrium identified in Proposition 1 is an ESS (for any n) and for all sufficiently small positive costs.*

Proof. By Proposition 1 the strategy profile is a symmetric Nash equilibrium and thus satisfies the first of the two conditions in Lemma 1.

For strictly positive costs pure strategy C yields a strictly lower payoff than the equilibrium payoff. However, all strategies $\sigma' \in \Delta(S)$ with $\sigma'(C) = 0$ satisfy $u(\sigma',\sigma) = u(\sigma,\sigma)$. We need to show that for all such $\sigma' \neq \sigma$ we have $u(\sigma,\sigma') - u(\sigma',\sigma') > 0$. As $\sigma'(C) = \sigma(C) = 0$ and using Equations 1 and 2, we obtain

$$\begin{array}{ll} u(\sigma,\sigma') - u(\sigma',\sigma') & = \\ (\sigma(A) - \sigma'(A)) \left(\alpha \frac{1-(1-\sigma'(A))^n}{\sigma'(A)} + (1-\alpha)\frac{1-\sigma'(B)^n}{1-\sigma'(B)}\right) & + \\ (\sigma(B) - \sigma'(B)) \left(\alpha\eta \frac{1-\sigma'(A)^n}{1-\sigma'(A)} + (1-\alpha)\lambda\frac{1-(1-\sigma'(B))^n}{\sigma'(B)}\right) & + \\ (\sigma(R) - \sigma'(R)) \left(\alpha\eta \frac{1-\sigma'(A)^n}{1-\sigma'(A)} + (1-\alpha)\frac{1-\sigma'(B)^n}{1-\sigma'(B)} - c\right). & \end{array}$$

As $\sigma(R) = 1 - \sigma(A) - \sigma(B)$ (and the same for σ'), this is equivalent to

$$\begin{array}{ll} u(\sigma,\sigma') - u(\sigma',\sigma') & = \\ \alpha\left(\sigma'(A) - \sigma(A)\right)\left(\frac{1-\sigma'(A)^n}{1-\sigma'(A)}\eta - \frac{1-(1-\sigma'(A))^n}{\sigma'(A)} - \frac{c}{\alpha}\right) & + \\ (1-\alpha)\left(\sigma'(B) - \sigma(B)\right)\left(\frac{1-\sigma'(B)^n}{1-\sigma'(B)} - \lambda\frac{1-(1-\sigma'(B))^n}{\sigma'(B)} - \frac{c}{1-\alpha}\right). & \end{array}$$

Note that last term in brackets in the first line is the equilibrium equation 3 (here for $\sigma'(A)$) and similarly the last term in brackets in the second line is the equilibrium equation 4 (here for $\sigma'(B)$). These terms are, therefore

zero, for $\sigma' = \sigma$ and they are, for all $n \geq 2$, increasing functions in $\sigma'(A)$ and $\sigma'(B)$, respectively. To see this for the first of these two expressions,

$$\frac{1-\sigma'(A)^n}{1-\sigma'(A)}\eta - \frac{1-(1-\sigma'(A))^n}{\sigma'(A)} - \frac{c}{\alpha}$$

note that the first term in this expression is

$$r_{B\eta} = \frac{1-\sigma'(A)^n}{1-\sigma'(A)} = \sum_{k=0}^{n-1} \sigma'(A)^k,$$

and, thus, clearly increasing in $\sigma'(A)$. Analogous arguments can be made for the second term in this first expression, as well as for both terms in the second of the two expressions. This means that the first term in brackets is negative exactly when $\sigma'(A) - \sigma(A)$ is negative and is positive exactly when the latter is positive. The same argument is true for the second term in brackets. Therefore, for any $\sigma' \neq \sigma$ with $\sigma'(C) = 0$ we, indeed, have $u(\sigma, \sigma') - u(\sigma', \sigma') > 0$. By Lemma 1 the equilibrium σ is, thus, an ESS. □

D Proof of Proposition 2

The proof proceeds as follows. That the symmetric Nash equilibrium identified in Proposition 1 is an ESS in the sense of Palm [1984] is shown in Proposition 4 in the previous section. The key step to prove the remaining claims in Proposition 2 is to prove that our game is a stable game, in the sense of Hofbauer and Sandholm [2009]. The remaining statements of Proposition 2 then follow from results in Hofbauer and Sandholm [2009].

In order to prove that our games are stable games, we need to introduce some of the more general notation introduced by Hofbauer and Sandholm [2009]. As we have just a single populations of players, in their notation we have $p = 1$. In their notation X denotes the space of players' mixed strategy space, which in our setting is, therefore, equal to $\Delta(S)$. The payoff function in their setting is denoted by $F : X \to \mathbf{R}^k$, where k is the number of pure strategies, i.e., $k = 4$ in our setting.[9] Translated to our notation, F

[9] Hofbauer and Sandholm [2009] actually use the letter n instead of k, but for our paper n is reserved to denote the number of players.

is, therefore, given by

$$F : \sigma \mapsto \begin{bmatrix} u(A,\sigma) \\ u(B,\sigma) \\ u(R,\sigma) \\ u(C,\sigma) \end{bmatrix}.$$

Definition 1 (Hofbauer and Sandholm [2009])**.** *A population game $F : X \to \mathbf{R}^k$ is a stable game if,*

$$(y-x)'(F(y)-F(x)) \leq 0, for\ all\ x, y \in X$$

We now prove that our games are stable games in this sense. We utilize [Hofbauer and Sandholm, 2009, Example 2.9] to show that our games are *negative dominant diagonal games*, which by [Hofbauer and Sandholm, 2009, Proposition 2.7] are stable games.

To do so we need to check two conditions. First, a condition of universal *negative frequency dependence*,

$$\frac{\partial u(s,\sigma)}{\partial \sigma(s)} \leq 0 \text{ for all } s \in \{A, B, R, C\},$$

and second,

$$\left|\frac{\partial u(s,\sigma)}{\partial \sigma(s)}\right| \geq \frac{1}{2}\left(\sum_{s' \neq s}\left|\frac{\partial u(s',\sigma)}{\partial \sigma(s)}\right| + \sum_{s' \neq s}\left|\frac{\partial u(s,\sigma)}{\partial \sigma(s')}\right|\right) \text{ for all } s \in \{A, B, R, C\}.$$

We here only provide the proof that these two conditions are satisfied for $s = A$, the proof for other strategies in $S = \{A, B, R, C\}$ proceeds analogously.

Recall that, by equation 1, using the notation from Appendix A,

$$\begin{bmatrix} u(A,\sigma) \\ u(B,\sigma) \\ u(R,\sigma) \\ u(C,\sigma) \end{bmatrix} = \begin{bmatrix} \alpha r_{A\eta} + (1-\alpha) r_{A\lambda} \\ \alpha r_{B\eta} + (1-\alpha) r_{B\lambda} \\ \alpha r_{B\eta} + (1-\alpha) r_{A\lambda} - c \\ \alpha r_{A\lambda} + (1-\alpha) r_{B\lambda} - c \end{bmatrix},$$

where, by Equation 2,

$$\begin{array}{rcl} r_{A\lambda} & = & \frac{1-(1-\sigma(A)-\sigma(R))^n}{\sigma(A)+\sigma(R)} \\ r_{A\eta} & = & \frac{1-(1-\sigma(A)-\sigma(C))^n}{\sigma(A)+\sigma(C)} \\ r_{B\lambda} & = & \lambda \cdot \frac{1-(1-\sigma(B)-\sigma(C))^n}{\sigma(B)+\sigma(C)} \\ r_{B\eta} & = & \eta \cdot \frac{1-(1-\sigma(B)-\sigma(R))^n}{\sigma(B)+\sigma(R)}. \end{array}$$

We, therefore, have that

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} = \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)} + (1-\alpha)\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)}.$$

Note that

$$\begin{array}{rcl} r_{A\lambda} & = & \frac{1-(1-\sigma(A)-\sigma(R))^n}{\sigma(A)+\sigma(R)} \\ & = & \sum_{j=0}^{n-1}(1-\sigma(A)-\sigma(R))^j, \end{array}$$

by the well-known formula for the geometric series. Thus, $\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)} \leq 0$. By the same argument $\frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)}$ is also everywhere non-positive, and, therefore,

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} \leq 0,$$

as desired. Analogous arguments yield that the desired first condition is satisfied:

$$\frac{\partial u(s,\sigma)}{\partial \sigma(s)} \leq 0 \text{ for all } s \in \{A,B,R,C\}.$$

Let us now consider the second condition, again for $s = A$, which can be written as

$$\begin{array}{rcl} \left|\frac{\partial u(A,\sigma)}{\partial \sigma(A)}\right| & \geq & \frac{1}{2}\left(\left|\frac{\partial u(B,\sigma)}{\partial \sigma(A)}\right| + \left|\frac{\partial u(R,\sigma)}{\partial \sigma(A)}\right| + \left|\frac{\partial u(C,\sigma)}{\partial \sigma(A)}\right|\right) \\ & + & \frac{1}{2}\left(\left|\frac{\partial u(A,\sigma)}{\partial \sigma(B)}\right| + \left|\frac{\partial u(A,\sigma)}{\partial \sigma(R)}\right| + \left|\frac{\partial u(A,\sigma)}{\partial \sigma(C)}\right|\right). \end{array}$$

We prove this statement in two steps. Note first that

$$u(A,\sigma) + u(B,\sigma) = u(R,\sigma) + u(C,\sigma) + 2c.$$

Therefore,

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} + \frac{\partial u(B,\sigma)}{\partial \sigma(A)} = \frac{\partial u(R,\sigma)}{\partial \sigma(A)} + \frac{\partial u(C,\sigma)}{\partial \sigma(A)}.$$

Since $\frac{\partial u(B,\sigma)}{\partial \sigma(A)} = 0$, we then have

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} = \frac{\partial u(B,\sigma)}{\partial \sigma(A)} + \frac{\partial u(R,\sigma)}{\partial \sigma(A)} + \frac{\partial u(C,\sigma)}{\partial \sigma(A)}.$$

As all these derivatives are either negative or equal to zero, we have

$$\left|\frac{\partial u(A,\sigma)}{\partial \sigma(A)}\right| = \left|\frac{\partial u(B,\sigma)}{\partial \sigma(A)}\right| + \left|\frac{\partial u(R,\sigma)}{\partial \sigma(A)}\right| + \left|\frac{\partial u(C,\sigma)}{\partial \sigma(A)}\right|.$$

For the second step, note that $\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)} = \frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(C)}$, as well as $\frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)} = \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(R)}$. Moreover, because $\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(R)} = \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(C)} = 0$, we have,

$$\begin{aligned}
\frac{\partial u(A,\sigma)}{\partial \sigma(A)} &= \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)} + (1-\alpha)\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)} \\
&= \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(R)} + (1-\alpha)\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(C)} + (1-\alpha)\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(R)} + \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(C)} \\
&= \frac{\partial u(A,\sigma)}{\partial \sigma(R)} + \frac{\partial u(A,\sigma)}{\partial \sigma(C)}.
\end{aligned}$$

As $\frac{\partial u(A,\sigma)}{\partial \sigma(B)} = 0$, and again, as all derivatives are either negative or equal to zero we have

$$\left|\frac{\partial u(A,\sigma)}{\partial \sigma(A)}\right| = \left|\frac{\partial u(A,\sigma)}{\partial \sigma(B)}\right| + \left|\frac{\partial u(A,\sigma)}{\partial \sigma(R)}\right| + \left|\frac{\partial u(C,\sigma)}{\partial \sigma(B)}\right|.$$

Finally, by adding the result of the two parts, the second condition follows for $s = A$. For strategies B, R, C the same result can be proven in an analogous manner.

Having, thus, established that we are dealing with stable games, the remaining results follow from results in Hofbauer and Sandholm [2009]: Theorem 5.1 for the BNN dynamics of Brown and Von Neumann [1950], Theorem 5.3 for the Best response dynamics of Gilboa and Matsui [1991], Theorem 7.1 for the Smith dynamics of Smith [1984], results from Section 7.2 for the replicator dynamics of Taylor and Jonker [1978] and the projection dynamics of Nagurney and Zhang [1997], and results from Section 4 for the Logit dynamics of Fudenberg and Levine [1998].

E Proof of Proposition 3

Let $u(A, G, x)$ and $u(B, G, x)$ denote the expected payoff to an individual who goes to food source A and B, respectively, when all others use mixed

strategy G and when the realized state of food availability at food source B is given by x. As $G(x) = P(y \leq x)$ is the probability that a random other individual goes to food source B, using again [Chao and Strawderman, 1972, Equation 3.4], we have

$$u(A, G, x) = \frac{1 - G(x)^n}{1 - G(x)}$$

and

$$u(B, G, x) = x \frac{1 - (1 - G(x))^n}{G(x)}.$$

The expected payoff for an individual using cut-off $y \in \chi$, when all others use mixed strategy G, is then given by

$$u(y, G) = \int_{x=x_L}^{y} \frac{1 - G(x)^n}{1 - G(x)} f(x)dx + \int_{x=y}^{x_H} x \frac{1 - (1 - G(x))^n}{G(x)} f(x)dx.$$

For a completely mixed strategy G to be a symmetric Nash equilibrium we need that $u(y, G)$ is constant for all $y \in \chi$. If $u(y, G)$ is a constant function in y then its derivative with respect to y must be constant as well (and equal to zero). This implies that

$$\frac{1 - G(y)^n}{1 - G(y)} f(y) - y \frac{1 - (1 - G(y))^n}{G(y)} f(y) = 0,$$

for all $y \in \chi$. Note that $f(y)$ cancels out and the resulting equilibrium distribution G does again not depend on the stochasticity of the food source, and we obtain

$$\frac{1 - G(y)^n}{1 - G(y)} = y \frac{1 - (1 - G(y))^n}{G(y)}.$$

Note that the left hand side of this equation is increasing in $G(y)$ and the right hand side is decreasing in $G(y)$. Note also that the left hand side is strictly less than the right hand side when both are evaluated at $G(y) = 0$ and the inequality is reversed when both are evaluated at $G(y) = 1$. Thus, for any $y \in \chi$, this equation has a unique solution $G(y)$. Note that this solution $G(y)$ must necessarily be higher for higher y. Note also that, necessarily, $G(0) = 0$ and $\lim_{y \to \infty} G(y) = 1$. G is, thus, a cumulative distribution function.

For finite n this solution is not easily expressed analytically. In the limit as $n \to \infty$ we obtain

$$\frac{1}{1 - G(y)} - y \frac{1}{G(y)} = 0,$$

for all $y \in \chi$, and, thus, $G(y) = \frac{y}{1+y}$ in this limiting case.

References

Alison M Bell, Shala J Hankison, and Kate L Laskowski. The repeatability of behaviour: a meta-analysis. *Animal Behaviour*, 77(4):771–783, 2009.

Ralph Bergmüller and Michael Taborsky. Animal personality due to social niche specialisation. *Trends in Ecology & Evolution*, 25(9):504–511, 2010.

Theodore C Bergstrom. On the evolution of hoarding, risk-taking, and wealth distribution in nonhuman and human populations. *Proceedings of the National Academy of Sciences*, 111(Supplement 3):10860–10867, 2014.

Daniel I Bolnick, Richard Svanbäck, James A Fordyce, Louie H Yang, Jeremy M Davis, C Darrin Hulsey, and Matthew L Forister. The ecology of individuals: incidence and implications of individual specialization. *The American Naturalist*, 161(1):1–28, 2003.

W Thomas Boyce. Differential susceptibility of the developing brain to contextual adversity and stress. *Neuropsychopharmacology*, 41(1):142–162, 2016.

M Broom, C Cannings, and GT Vickers. Multi-player matrix games. *Bulletin of Mathematical Biology*, 59(5):931–952, 1997.

George W Brown and John Von Neumann. Solutions of games by differential equations. Technical report, Rand Corporation, Santa Monica, CA, 1950.

Min-Te Chao and WE Strawderman. Negative moments of positive random variables. *Journal of the American Statistical Association*, 67(338):429–431, 1972.

Sasha RX Dall, Alison M Bell, Daniel I Bolnick, and Francis LW Ratnieks. An evolutionary ecology of individual differences. *Ecology Letters*, 15(10): 1189–1198, 2012.

Niels J Dingemanse and Max Wolf. Recent models for adaptive personality differences: a review. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1560):3947–3958, 2010.

Bruce J Ellis and W Thomas Boyce. Biological sensitivity to context. *Current Directions in Psychological Science*, 17(3):183–187, 2008.

Bruce J Ellis, W Thomas Boyce, Jay Belsky, Marian J Bakermans-Kranenburg, and Marinus H Van Ijzendoorn. Differential susceptibility to the environment: An evolutionary–neurodevelopmental theory. *Development and Psychopathology*, 23(1):7–28, 2011.

Jeffrey C Ely and William H Sandholm. Evolution in bayesian games i: Theory. *Games and Economic Behavior*, 53(1):83–109, 2005.

Willem E Frankenhuis, Karthik Panchanathan, and Jay Belsky. A mathematical model of the evolution of individual differences in developmental plasticity arising through parental bet-hedging. *Developmental Science*, 19(2):251–274, 2016.

Stephen Dewitt Fretwell. On territorial behavior and other factors influencing habitat distribution in birds. *Acta Biotheoretica*, 19(1):45–52, 1969.

Drew Fudenberg and David K Levine. *The Theory of Learning in Games*. MIT press, 1998.

Itzhak Gilboa and Akihiko Matsui. Social stability and equilibrium. *Econometrica*, 59:859–867, 1991.

Stephanie M Harris, Sébastien Descamps, Lynne U Sneddon, Philip Bertrand, Olivier Chastel, and Samantha C Patrick. Personality predicts foraging site fidelity and trip repeatability in a marine predator. *Journal of Animal Ecology*, 89(1):68–79, 2020.

John C Harsanyi. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1):1–23, 1973.

Josef Hofbauer and William H Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 144(4):1665–1693, 2009.

David J Jörg, Yu Kitadate, Shosei Yoshida, and Benjamin D Simons. Competition for stem cell fate determinants as a mechanism for tissue homeostasis. arXiv preprint arXiv:1901.03903, 2019.

Elon Kohlberg and Jean-François Mertens. On the strategic stability of equilibria. *Econometrica*, 54(5):1003–37, 1986.

Vlastimil Křivan, Ross Cressman, and Candace Schneider. The ideal free distribution: a review and synthesis of the game-theoretic perspective. *Theoretical Population Biology*, 73(3):403–425, 2008.

Francesca Lionetti, Arthur Aron, Elaine N Aron, G Leonard Burns, Jadzia Jagiellowicz, and Michael Pluess. Dandelions, tulips and orchids: Evidence for the existence of low-sensitive, medium-sensitive and high-sensitive individuals. *Translational Psychiatry*, 8(1):1–11, 2018.

John M McNamara and Alasdair I Houston. If animals know their own fighting ability, the evolutionarily stable level of fighting is reduced. *Journal of Theoretical Biology*, 232(1):1–6, 2005.

David J Mitchell and Thomas M Houslay. Context-dependent trait covariances: how plasticity shapes behavioral syndromes. *Behavioral Ecology*, 32(1):25–29, 2021.

Pierre-Olivier Montiglio, Caterina Ferrari, and Denis Reale. Social niche specialization under constraints: personality, social interactions and environmental heterogeneity. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1618):20120343, 2013.

John H Nachbar. "Evolutionary" selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19(1):59–89, 1990.

Anna Nagurney and Ding Zhang. Projected dynamical systems in the formulation, stability analysis, and computation of fixed-demand traffic network equilibria. *Transportation Science*, 31(2):147–158, 1997.

John F Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.

Daniel Nettle and Lars Penke. Personality: bridging the literatures from human psychology and behavioural ecology. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1560):4043–4050, 2010.

Petri T Niemelä and Niels J Dingemanse. Meta-analysis reveals weak associations between intrinsic state and personality. *Proceedings of the Royal Society B: Biological Sciences*, 285(1873):20172823, 2018.

Günther Palm. Evolutionary stable strategies and game dynamics for n-person games. *Journal of Mathematical Biology*, 19(3):329–334, 1984.

Michael Pluess. Individual differences in environmental sensitivity. *Child Development Perspectives*, 9(3):138–143, 2015.

William H Sandholm. Evolution in bayesian games ii: Stability of purified equilibria. *Journal of Economic Theory*, 136(1):641–667, 2007.

Andrew Sih, Alison Bell, and J Chadwick Johnson. Behavioral syndromes: an ecological and evolutionary overview. *Trends in Ecology & Evolution*, 19(7):372–378, 2004.

John M Smith and George R Price. The logic of animal conflict. *Nature*, 246(5427):15–18, 1973.

Michael J Smith. The stability of a dynamic model of traffic assignment—an application of a method of Lyapunov. *Transportation Science*, 18(3):245–252, 1984.

Orr Spiegel, Stephan T Leu, C Michael Bull, and Andrew Sih. What's your move? Movement as a link between personality and spatial dynamics in animal populations. *Ecology Letters*, 20(1):3–18, 2017.

Peter D Taylor and Leo B Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.

Benjamin J Toscano, Natasha J Gownaris, Sarah M Heerhartz, and Cristián J Monaco. Personality, foraging behavior and specialization: integrating behavioral and food web ecology at the individual level. *Oecologia*, 182(1):55–69, 2016.

Max Wolf and John M McNamara. On the evolution of personalities via frequency-dependent selection. *The American Naturalist*, 179(6):679–692, 2012.

Max Wolf and Franz J Weissing. An explanatory framework for adaptive personality differences. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1560):3959–3968, 2010.

Max Wolf, G Sander Van Doorn, Olof Leimar, and Franz J Weissing. Life-history trade-offs favour the evolution of animal personalities. *Nature*, 447 (7144):581–584, 2007.

Max Wolf, G Sander Van Doorn, and Franz J Weissing. Evolutionary emergence of responsive and unresponsive personalities. *Proceedings of the National Academy of Sciences*, 105(41):15825–15830, 2008.

Max Wolf, G Sander Van Doorn, and Franz J Weissing. On the coevolution of social responsiveness and behavioural consistency. *Proceedings of the Royal Society B: Biological Sciences*, 278(1704):440–448, 2011.