



ELSEVIER

Contents lists available at ScienceDirect

Journal of Theoretical Biology

journal homepage: www.elsevier.com/locate/yjtbi

Control in mutualisms: Combined implications of partner choice and bargaining roles

Antonio J. Golubski^{a,*}, Christopher A. Klausmeier^{a,b}

^a Michigan State University, W.K. Kellogg Biological Station, USA

^b Michigan State University, Department of Plant Biology, USA

ARTICLE INFO

Article history:

Received 18 January 2010

Received in revised form

13 September 2010

Accepted 15 September 2010

Available online 22 September 2010

Keywords:

Mutualism

Partner choice

Bargaining games

Cheater

Competition-colonization tradeoff

ABSTRACT

When two species form a mutualistic association, the degree of control that each has over the interaction may be pivotal in determining the relative benefit each obtains. We incorporate the capacity for partner choice into a model of mutualism based on the exchange of goods and/or services, where one guild of mutualists plays the role of proposer (proposing a price at which the goods and/or services will be exchanged) and the other plays the role of responder (accepting or rejecting the deal). We show how the payoff structure in this scenario and other closely related ones correspond to the ultimatum and demand games of economics. In the model, there are both costs and benefits to a guild whose players have control over interactions. Control over interactions in the sense of being able to exercise partner choice can benefit a guild by selecting for mutualism in its partners, but is most effective in selecting against moderately exploitative partners, and so can give highly exploitative partners an advantage. This can generate dynamics similar to taxon cycles or those seen in models with competition-colonization tradeoffs, wherein increasingly more mutualistic partners (acting as superior competitors) are selected for up to a tipping point, at which highly exploitative strategies (akin to superior colonizers) gain an advantage. Control over interactions in the sense of being able to appropriate 'surplus' payoffs in each interaction, which is selected for within-guild and is equivalent to playing the role of responders, selects against high demands (and so for mutualism) in the guild with control. Combining the two mechanisms, a high degree of mutualism in both guilds and coexistence of more mutualistic and more exploitative strategies within each are both consistent with control over the interaction being highly skewed toward one side that does what is in its own short-term interests.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

It is well accepted that when two mutualists interact, their interests often will not be perfectly aligned (e.g. Herre et al., 1999). For example, the carbon that a plant or aphid provides an ant protector, mycorrhizal fungus, or pollinator is presumably often costly for the giver to produce and valuable to the receiver. Thus the donor would benefit if less carbon were exchanged for each unit of good or service its partner provides, while the recipient would benefit if more carbon were exchanged. This conflict of interests introduces a zero-sum aspect to mutualisms that sets up a temptation to 'cheat' partners by becoming less mutualistic and more exploitative towards them.

When mutualism is modeled as trade, increasing the payoff to one participant in an interaction generally comes at the expense of decreasing the payoff to its partner (as for example increasing

the price paid for a good or service benefits the provider at the expense of the buyer, in line with the examples of conflicts of interest given earlier). Highly mutualistic species are thus those species that interact on terms highly beneficial to their partners (e.g. Noë et al., 1991; Schwartz and Hoeksema, 1998; Bergstrom and Lachmann, 2002; Johnstone and Bshary, 2002, 2007, 2008; Hoeksema and Schwartz, 2003; Akçay and Roughgarden, 2007; Friesen and Mathias, 2010). We adopt that perspective here, and define mutualism as willingness to enter into associations with terms relatively more favorable to one's partner (i.e. demanding only a relatively low payoff from the association).

The degree of control that each partner exercises over an interaction may be central to explaining both how the mutualistic relationship persists in the face of potential cheaters/exploiters, and the relative degree of benefit each partner obtains. Preferential association with or rewarding of partners that are more mutualistic (partner choice *sensu* Bull and Rice, 1991), is a key way in which mutualists may exert control over an interaction, and is one of the most common mechanisms offered to explain the persistence of mutualism in the face of cheating (e.g. Noë and Hammerstein, 1994, 1995; Denison, 2000; Noë, 2001; Simms and Taylor, 2002; West et al., 2002; Hoeksema and Kummel, 2003;

* Corresponding author. Current address: University of Toronto, Department of Ecology and Evolutionary Biology, 25 Harbord Street, Toronto, ON, Canada M5S 3G5. Tel.: +1 416 946 7223.

E-mail addresses: antonio.golubski@utoronto.ca (A.J. Golubski), klausme1@msu.edu (C.A. Klausmeier).

Sachs et al., 2004; Foster and Wenseleers, 2006; West et al., 2007; Friesen and Mathias, 2010). Empirical work has found evidence suggesting partner choice by mutualists as diverse as ant-plants (Edwards et al., 2006), cleaner fish clients (Bshary and Grutter, 2002; Bshary and Schäffer, 2002), and plants in mycorrhizal (Bever et al., 2009) and rhizobial (Kiers et al., 2003; Simms et al., 2006) symbioses. However, if partner choice is responsible for maintaining mutualistic partners, guilds unable to exercise such choice might be expected to suffer high degrees of exploitation, begging the question as to what keeps the association mutualistic for them. Recent models have shown that a mutually beneficial association can be maintained if both partner guilds engage in choice (Ferriere et al., 2002, 2007) or if members of one guild are able to exploit partners while members of the opposite guild can determine whether and/or for how long an interaction occurs (Johnstone and Bshary, 2002, 2007, 2008), but it remains unclear what mechanism(s) might ensure mutualism toward a guild that has little control over the interaction.

If the exchange of goods and services can occur on a range of different terms while remaining profitable for both parties, potential 'surplus' payoffs may exist beyond the minimum required by each party to participate in the association. Economic game theory suggests the interesting possibility that in this situation a guild with the upper hand in the short-term (the one obtaining any 'surplus' in each interaction), might lose out over the long term because of selective pressures on the payoff it demands (Gale et al., 1995; Toquenaga and Suzuki, 2005; Poulsen, 2007). This yields a very different outcome than the fairness that is often expected based on the Nash bargaining solution (Nash, 1950). It also provides a strikingly different mechanism by which mutualistic partners might be selected for than does partner choice: a sort of indiscriminate generosity (in the form of giving away of surpluses) rather than targeted sanctions (in the form of shunning exploitative partners that demand high payoffs). Conversely, it presents a scenario in which having control of an interaction (in the sense of having the power to appropriate the surplus) may be disadvantageous. What then are the implications of control over a mutualistic interaction when some mechanisms may cause that control to be advantageous while other mechanisms cause it to be disadvantageous? To address that question, we develop a trade-based model of mutualism in which parties differ in both their bargaining roles and their capacity for partner choice, and examine the selective pressures on strategies ranging from more mutualistic to more exploitative in each guild.

2. Models and results

We model a scenario in which mutualistic associations form between pairs of individuals from two guilds (*A* and *B*), representing two species of mutualist. Let D_1, D_2, \dots, D_n represent the range of all possible terms on which mutualistic associations may occur (what we will call 'deals'), ranked by their relative profitability to a player of guild *B*. Thus, if we let $b(D_x)$ give the fitness that a guild *B* player achieves (the player's 'payoff') when it forms an association on terms D_x , $b(D_1) < b(D_2) < \dots < b(D_n)$. Values of x may be discrete or continuous, so that the set of possible deals may be finite or infinite. However, because many of our results are based on considering slight payoff differences, they may not apply if possible deals are too coarse-grained. These deal profitabilities may involve prices at which goods and/or services are exchanged as well as volumes of goods and services traded, or any other terms of an association that affect that association's profitability for a guild *B* player. Let the ranks of deals' profitabilities to guild *A* players be inverted relative to their profitabilities to guild *B* players, so that if $a(D_x)$ gives the profitability of deal D_x to a guild

A player, $a(D_1) > a(D_2) > \dots > a(D_n)$. Thus a conflict of interests exists so that the more favorable a deal is for one partner, the less favorable it is for the other.

Let α denote a guild *A* strategy and equal the lowest profitability (highest rank) deal that a guild *A* player with that strategy is willing to accept. Thus a guild *A* player with strategy α demands a fitness payoff of at least $a(D_\alpha)$. Let β denote a guild *B* strategy and equal the lowest profitability (lowest rank) deal that a guild *B* player with that strategy will accept, so that the guild *B* player demands a payoff of at least $b(D_\beta)$. When two players interact and $\alpha > \beta$, there exists a distance in terms of deal rank space of $\alpha - \beta$ that can allow one or both players to achieve excess payoffs beyond the minimum dictated by their demands; we will refer to this as the 'surplus'. For example, if a guild *A* player with strategy $\alpha = 7$ (which will therefore accept any payoff $a(D_1)$ through $a(D_7)$) and a guild *B* player with strategy $\beta = 3$ (which will accept any payoff $b(D_3)$ through $b(D_n)$) interact, the surplus is 4. It is important to note that we do not specify what this 'surplus' or any fraction of it translates to in terms of potential payoffs for each player. The surplus or any portion of it may be highly valuable to one player and of little value to the other, or nonlinearities in the relationships between deals and payoffs may mean that the value of a given fraction of a given surplus to a player depends sensitively on both partners' specific demands. Let s give the fraction of the surplus that benefits the guild *A* player and $(1 - s)$ give the fraction that benefits the guild *B* player, so that when a successful association occurs between two players the final deal is $D_{\alpha - s(\alpha - \beta)} = D_{\beta + (1 - s)(\alpha - \beta)}$. We assume that s is fixed, so that the fraction of the surplus obtained is identical for all members of a guild and does not change over time. Thus in the above example where $\alpha = 7$ and $\beta = 3$, payoffs to each player when $s = 0.5$ will be $a(D_{7 - 0.5(7 - 3)}) = a(D_5)$ and $b(D_{3 + (1 - 0.5)(7 - 3)}) = b(D_5)$. Let ω_A and ω_B give the payoffs to guild *A* and *B* players, respectively, when they do not form associations. We assume that ω_A and ω_B are non-negative, independent of player strategies, and lower than the minimum payoff received from any successful association (equivalent to assuming that players never demand so little payoff from an association that they would achieve higher payoffs on their own). A minimum payoff of zero is consistent with an obligate mutualism.

2.1. No partner choice

We first consider the case where there is no partner choice; each player in guild *A* interacts with one player in guild *B*. In this case the payoffs to guild *A* and *B* players with strategies α and β , respectively, are:

$$p_\alpha = \begin{cases} a(D_{\beta + (1-s)(\alpha-\beta)}) & \text{when } \beta \leq \alpha, \\ \omega_A & \text{otherwise} \end{cases} \quad (1)$$

and

$$p_\beta = \begin{cases} b(D_{\beta + (1-s)(\alpha-\beta)}) & \text{when } \beta \leq \alpha, \\ \omega_B & \text{otherwise} \end{cases} \quad (2)$$

We focus on the case where one guild receives the entire surplus, and let $s = 1$ so that this is guild *A*. This is consistent with guild *A* players responding to offers made by guild *B* players. When $s = 1$ and $\beta \leq \alpha$, Eq. (1) reduces to $a(D_\beta)$, meaning that when interacting with any strategy β , the guild *A* player's payoff is maximized purely by minimizing the probability of β being greater than α , which is accomplished by $\alpha = n$. Thus, the optimum strategy for guild *A* is to accept all guild *B* offers, including those that lead to the guild *A* player obtaining the minimum possible payoff from associations. For any distribution of α and β values,

the highest α (lowest demand) will achieve the maximum expected per capita payoff, and will therefore be expected to continually increase in frequency. Note however that if the highest β strategies are absent or die out, a range of high α strategies will all achieve the same maximum payoff (for example if the highest β present is 5, any $\alpha \geq 5$ will obtain an equivalent payoff).

When $s=1$ and $\beta \leq \alpha$, Eq. (2) reduces to $b(D_\beta)$, so the only factor preventing $\beta=n$ (demanding the highest possible payoff) from being the ideal strategy for members of guild B is the imperative to keep β less than or equal to α . However, since from Eq. (1) α is expected to approach n , β approaching n will ultimately become the optimum strategy. Thus, while obtaining the surplus is beneficial to members of guild A over the short term, it exerts selective pressure for low demands (Table 1), which over evolutionary time favors the opposite guild. This result has been shown previously in similar models (Güth et al., 1982; Gale et al., 1995; Toquenaga and Suzuki, 2005; Fischer et al., 2006; Poulsen, 2007). The specific result of β being able to approach n depends on some frequency of high β strategies being present to drive α towards n , and thus is most appropriate for scenarios where immigration, mutation, or behavioral plasticity keeps all possible strategies present at some small frequency.

We do not model variability between competing strategies within a guild in the amount of the surplus they keep. However, because payoffs increase with s in Eq. (1) and decrease with s in Eq. (2), it is straightforward that within a guild, strategies that keep more of the surplus would outperform strategies that keep less (as in Gale et al., 1995; Toquenaga and Suzuki, 2005; Poulsen, 2007). Thus, the value taken by the fixed s we assume would reflect details of the interaction between partners that determine the fraction of the surplus that players from each guild are ultimately able to appropriate.

Payoffs in our model without partner choice and when $s=1$ correspond to the demand and ultimatum games in economics (Güth et al., 1982; Gale et al., 1995; Toquenaga and Suzuki, 2005; Fischer et al., 2006; Poulsen, 2007). In the demand game, both players demand fractions of a resource to be split, and players on one side may preferentially receive surpluses remaining after both parties' demands are satisfied (although surpluses are also

often assumed to go unclaimed by either player, a scenario our model cannot reproduce). In the ultimatum game, 'proposers' demand a fraction of a resource and 'responders' accept or reject those offers, which equates to responders (which receive the remaining fraction of the resource) receiving surpluses. Responders have control in the sense of determining whether or not interactions occur, and their payoffs correspond to players in the demand game which obtain the entire surplus, and thus profit maximally from each interaction. As illustrated above, these individual-level advantages have costs in the form of selecting for low demands in the guild as a whole. Similar to the results of recent work bridging the demand and ultimatum games (Fischer et al., 2006), simulations of our model with intermediate s values showed increasing fairness as s approached 0.5 (supplementary materials figure S4).

Congruence between the above payoffs and the demand and ultimatum games does not imply that our model assumes that mutualists divide a resource, only that when an association occurs both parties' payoffs are functions of a 'deal' that can be determined by the strategy of one partner. To illustrate and elaborate on this, consider the cartoonish example of customers shopping for candy bars. If we assume that stores will only sell bars at prices that are profitable, the minimum payoff for a store will occur when no bars are sold (because our model considers one association per round, it only compares the sale of a single candy bar at various prices with no sale). A store that sets the price for a candy bar at \$1 determines the net payoff to a customer: the joy that customer gains from eating the candy bar, minus \$1. This payoff is the same for all customers whether they would have been willing to pay \$1, \$2, or \$5, and so if customers represent guild A, the situation is consistent with $s=1$. Assume that eating a candy bar is worth enough to a customer that all three prices represent good deals, with a cheaper price merely representing a better deal. The $s=0$ case would be equivalent to customers paying the maximum price they would be willing to pay (\$1, \$2, or \$5) regardless of how much lower the price offered by the store is. The $s=0.5$ case would have the three customers described above paying \$1, \$1.50, and \$3, respectively: halfway between their threshold and the price the store demands. The fact that the behaviors described in the $s=0$ and 0.5 cases are rarely seen reflects the fact that obtaining surpluses is beneficial at the level of the individual. The absence of this behavior from our everyday experience may also contribute to making less intuitively obvious the selective pressure that it would exert on stores to lower their prices. Extending the above example to mutualisms only requires that paying lower prices is preferable to paying higher prices, so that like the dollar amounts above, the currency in which the price is paid is valuable or costly to the party paying it. This will be true in mutualisms whenever partners could profit from allocating the time and/or energy invested into providing mutualistic benefits elsewhere. This is reasonable for example in mycorrhizal mutualisms, where plant growth requires both soil nutrients (which plants can often obtain more effectively through trade with mycorrhizal fungi than with their own roots) and carbon (the currency in which those fungi are paid). Rhizobial symbioses present a similar situation. It is reasonable to suspect that organisms which invest resources into creating structures or compounds that attract or reward ant protectors, pollinators, or seed dispersers could often benefit from being able to put those resources to other uses if doing so would not cause a reduction in benefits provided by their mutualistic partners. Bites of mucus or healthy tissue that cleaner fish may demand as a price for service above what they obtain from parasites eaten are also presumably costly for client fish.

Next we consider the implications of customers shopping around across multiple candy shops.

Table 1

Payoffs to various guild A strategies when paired with various guild B strategies, when $s=1$ (upper triangles) versus $s=0$ (lower triangles). We assume here that the surplus gained equals one minus the sum of the demands all multiplied by s for the guild A player or $(1-s)$ for the guild B player, and that demands must sum to ≤ 1 for a successful association to occur. An "X" indicates where a successful association did not occur. Bottom row: average payoffs to each guild A strategy when the guild B population consists of strategies 0.2, 0.4, 0.6, and 0.8 in frequencies 0.2, 0.5, 0.2, and 0.1. For any pairing of strategies, the guild A player obtains as high or higher payoffs when $s=1$ than when $s=0$, but when $s=1$ the most mutualistic guild A strategy is always expected to achieve the highest payoffs, which need not be the case when $s=0$.

		Guild A player demand				
		0.9	0.7	0.5	0.3	0.1
Guild B player demand	0.8	X	X	X	X	0.2
	0.6	X	X	X	0.4	0.4
	0.4	X	X	0.6	0.6	0.6
	0.2	X	0.8	0.8	0.8	0.8
	Average	X	0.16	0.46	0.54	0.56
		$(s=0)$				
		X	0.14	0.35	0.27	0.10

2.2. Partner choice

We introduce partner choice by allowing members of one guild to each assess two individuals of the opposite guild and associate with only the more mutualistic of the two. If both of the latter play the same strategy, one is associated with at random. In mycorrhizal mutualisms for example, this would be consistent with each plant having two root segments that each became infected by a single fungal partner, and then providing carbon to (and receiving nutrients from) only the root segment hosting the more mutualistic fungus. It would be consistent with similar scenarios occurring across multiple nodules in rhizobial mutualisms or across multiple domatia in some ant-plant mutualisms. It would also be consistent with pollinators or seed dispersers allocating their time and effort to only the most rewarding of multiple potential plant partners. This scenario would also be consistent with a habitat that is more finely divided for the guild being subjected to choice (e.g. because they are smaller or less mobile), or with a relative shortage of partners in the choosing guild and abundance of partners in the chosen guild. For simplicity, we will say that members of the choosing guild occupy 'sites', and at each site are two 'slots' that may be occupied by members of the guild experiencing choice. We later consider the possibility that slots may occasionally be empty (e.g. because of random disturbances). When this happens, the player in the choosing guild will associate with the remaining partner, so long as their demands are compatible. Subscripts $f, g, h, i, j,$ and k denote different strategies within a guild, and are used only when such distinctions are necessary. The relative value of subscripts are not implied by their alphabetical order, and are only assumed where indicated in equations or in the text. Note that based on our notation high α values indicate more mutualistic guild A strategies, while more mutualistic guild B strategies are indicated by low β values.

For simplicity, we assume that $s=1$ for Eqs. (3)–(10). First let us give the capacity for partner choice to guild B, which does not receive the surplus. If a member of guild B playing strategy β is able to choose to associate with the more mutualistic of two members of guild A, playing $\alpha=i$ and $\alpha=j$, and assuming $i \neq j$, the payoff to $\alpha=i$ that follows from Eqs. (1) and (2) is

$$p_{\alpha=i} = \begin{cases} a(D_\beta) & \text{when } \beta \leq (\alpha=i) \text{ and } (\alpha=i) = \max[(\alpha=i), (\alpha=j)], \\ \omega_A & \text{otherwise,} \end{cases} \quad (3)$$

with the player having a 50% probability of getting each payoff listed above when $i=j$, and the same expression, with i and j reversed, giving the payoff to $(\alpha=j)$. The payoff to β is

$$p_\beta = \begin{cases} b(D_\beta) & \text{when } \beta \leq \max[(\alpha=i), (\alpha=j)], \\ \omega_B & \text{otherwise} \end{cases} \quad (4)$$

In this case partner choice does not change the qualitative outcome. The payoff to guild A is again maximized by α approaching n (demanding minimum payoff), and β will again be free to approach n (demanding maximum payoff).

More interesting results arise when we instead give partner choice to the guild receiving the surplus (which also seems the more realistic scenario, since it would correspond to one guild's players paying the lower of two competing prices). The payoff to a member of guild A playing strategy α when it can choose to interact with the more mutualistic of two partners (playing strategies $\beta=i$ and $\beta=j$), is

$$p_\alpha = \begin{cases} \max[a(D_i), a(D_j)] & \text{when } \alpha \geq \min[(\beta=i), (\beta=j)], \\ \omega_A & \text{otherwise} \end{cases} \quad (5)$$

The payoff to the member of guild B playing strategy $\beta=i$ (when $i \neq j$) is

$$p_{\beta=i} = \begin{cases} b(D_i) & \text{when } (\beta=i) \leq \alpha \text{ and } (\beta=i) = \min[(\beta=i), (\beta=j)], \\ \omega_B & \text{otherwise} \end{cases} \quad (6)$$

The player again has a 50% chance of receiving each payoff should $i=j$, and Eq. (6) with i and j reversed again gives the payoff to $\beta=j$. The final condition in the top row of Eq. (6) now provides a selective pressure in favor of lower guild B demands (lower β) even if all α approach n , which acts counter to the first term that favors higher β values.

2.3. Relative fitness of mutualists under choice

To explore the consequences of these opposing pressures when $s=1$, let us consider the relative fitnesses of two guild B strategies, $\beta=i$ and $\beta=j$, where $\beta=i$ is the more mutualistic of the two (i.e. $i < j$), and when only those two strategies are present (for now we will not consider empty slots). Let $F_{\beta=i}$ and $F_{\beta=j}$ be the frequencies of each guild B strategy, and let $F_{\alpha \geq i}$ and $F_{\alpha \geq j}$ be the frequency of guild A players willing to make associations with each (note that $F_{\alpha \geq i} > F_{\alpha \geq j}$). The expected payoff to each strategy per slot it occupies equals its demand, multiplied by the probability that its partner's demand is low enough for an association to occur, multiplied by its probability of being the most mutualistic of the two associations available to its partner (plus one half the probability of both the associations available to a partner being the focal strategy). To this is added the probability of failing to form an association, multiplied by the minimum payoff that occurs when no association is formed. This yields expected payoffs of

$$p_{\beta=i} = b(D_i)F_{\alpha \geq i} \left(\frac{1}{2}F_{\beta=i} + F_{\beta=j} \right) + \omega_B F_{\alpha \geq i} \frac{1}{2}F_{\beta=i} + \omega_B(1 - F_{\alpha \geq i}) \quad (7)$$

and

$$p_{\beta=j} = b(D_j)F_{\alpha \geq j} \left(\frac{1}{2}F_{\beta=j} \right) + \omega_B F_{\alpha \geq j} \left(\frac{1}{2}F_{\beta=j} + F_{\beta=i} \right) + \omega_B(1 - F_{\alpha \geq j}). \quad (8)$$

The condition for $\beta=j$ to increase in frequency (i.e. have higher per capita fitness than $\beta=i$) is that Eq. (8) be greater than Eq. (7), which reduces to

$$\frac{b(D_j) - \omega_B}{b(D_i) - \omega_B} > \left(\frac{F_{\alpha \geq i}}{F_{\alpha \geq j}} \right) \left[\frac{F_{\beta=i}}{F_{\beta=j}} + 2 \right]. \quad (9)$$

From Eq. (9), we can see that if two strategies have similar demands (the left hand side of the inequality is close to one), the more mutualistic strategy will be favored. When $\omega=0$ (i.e. an obligate mutualism), a more exploitative strategy can only do better than a more mutualistic competitor if it demands at least twice the payoff that the competitor does, and the rarer the less mutualistic strategy is, the higher its demand needs to be. The reason for the effect of higher demand is most clearly illustrated by assuming all members of guild A are willing to make any deal (and $F_{\alpha \geq i} = F_{\alpha \geq j} = 1$). In this case the costs of a higher demand (imposed by guild A's preferential association with the more mutualistic competitor) do not increase with the actual value of the demand, while the benefits of the higher demand (a higher payoff when an association does occur) do. The result is a situation where one competitor can gain a fitness advantage over

another either by being slightly more mutualistic or by being much less mutualistic, but does worst when it is slightly less mutualistic (Fig. 1a). The ω_B term facilitates the more exploitative $\beta=j$; the closer ω_B is to $b(D_i)$, the smaller the difference between $b(D_j)$ and $b(D_i)$ needed for $\beta=j$ to successfully invade. Additional guild A strategies can mitigate the advantage of very high demands, but do not change the fact that a slightly more demanding strategy does worse than a slightly less demanding one. Eq. (9) also shows that the less mutualistic strategy cannot invade when arbitrarily rare, no matter what its demand or its competitor's demand (Fig. 1b). The less mutualistic strategy relies on the payoff it receives from members of guild A that have no other choice but to associate with in (i.e. that encounter two partners both playing $\beta=j$), and the frequency of that scenario declines exponentially with the frequency of the strategy.

Considering the relative fitnesses of a more demanding guild B strategy $\beta=j$ and a less demanding strategy $\beta=i$ in the presence of any number of other strategies in both guilds, the condition for $\beta=j$ to increase in frequency relative to $\beta=i$ is

$$\frac{b(D_j)}{b(D_i)} > \left(\frac{F_{\alpha \geq i}}{F_{\alpha \geq j}} \right) \left[1 + \frac{(1 - (\omega_B/b(D_i)) \cdot (F_{\alpha \geq j}/F_{\alpha \geq i})) \left((1/2)F_{\beta=j} + (1/2)F_{\beta=i} + \sum_{i < h < j} F_{\beta=h} \right) + (\omega_B/b(D_i)) (1 - (F_{\alpha \geq j}/F_{\alpha \geq i})) \left((1/2)F_{\beta=i} + \sum_{g < i} F_{\beta=g} - 1 \right)}{(1/2)F_{\beta=j} + \sum_{j < k} F_{\beta=k}} \right]. \tag{10a}$$

The implications of Eq. (10) can be made more obvious by one of two alternate simplifying assumptions. When $\omega_B=0$ (an obligate mutualism), Eq. (10) reduces to

$$\frac{b(D_j)}{b(D_i)} > \left(\frac{F_{\alpha \geq i}}{F_{\alpha \geq j}} \right) \left[1 + \frac{(1/2)F_{\beta=j} + (1/2)F_{\beta=i} + \sum_{i < h < j} F_{\beta=h}}{(1/2)F_{\beta=j} + \sum_{j < k} F_{\beta=k}} \right]. \tag{10b}$$

When $(F_{\alpha \geq i} \approx F_{\alpha \geq j})$, it reduces to approximately

$$\frac{b(D_j)}{b(D_i)} > \left[1 + \frac{(1 - (\omega_B/b(D_i))) \left((1/2)F_{\beta=j} + (1/2)F_{\beta=i} + \sum_{i < h < j} F_{\beta=h} \right)}{(1/2)F_{\beta=j} + \sum_{j < k} F_{\beta=k}} \right]. \tag{10c}$$

Here, it is possible for $\beta=j$ to have only slightly higher demand than $\beta=i$ and still have higher per capita fitness, even when arbitrarily rare. This requires that: (1) most guild A members willing to associate with $\beta=i$ are willing to associate with $\beta=j$ ($F_{\alpha \geq i} \approx F_{\alpha \geq j}$), and (2) guild B strategies less mutualistic than $\beta=j$ are sufficiently common. Note also that in this case the less mutualistic $\beta=j$ can increase from arbitrarily low frequency, provided that even more exploitative strategies are sufficiently common. Because any association is preferable to no association at all for members of guild A, any frequency of empty slots would be added to the $F_{\beta=k}$ term, effectively acting as highly demanding strategies and benefiting $\beta=j$ over the more mutualistic $\beta=i$. In Eq. (10c), the ω_B term again facilitates the more exploitative $\beta=j$, although ω_B can also facilitate the more mutualistic $\beta=i$ in Eq. (10) when strategies more mutualistic than $\beta=i$ are sufficiently common.

General equations for the expected per capita payoffs to any strategies $\alpha=f$ and $\beta=i$ in guilds A and B, respectively, where there are z slots within which members of guild B occur for each member of guild A to choose among, and again allowing s to take

values other than 1, are

$$P_{\alpha=f} = \sum_{i \leq f} F_{\beta=i} a(D_{f-s(f-i)}) \left\{ 1 - \sum_{h < i} F_{\beta=h} \right\} + \omega_A \left[1 - \sum_{i \leq f} F_{\beta=i} \left\{ 1 - \sum_{h < i} F_{\beta=h} \right\} \right] \tag{11}$$

and

$$P_{\beta=i} = \sum_{i \leq f} F_{\alpha=f} b(D_{i+(1-s)(f-i)}) \left\{ \sum_{\psi=0}^{z-1} \frac{1}{\psi+1} \binom{z-1}{\psi} F_{\beta=i}^{\psi} \left(\sum_{k>i} F_{\beta=k} \right)^{z-1-\psi} \right\} + \omega_B \left[1 - \left(\sum_{i \leq f} F_{\alpha=f} \left\{ \sum_{\psi=0}^{z-1} \frac{1}{\psi+1} \binom{z-1}{\psi} F_{\beta=i}^{\psi} \left(\sum_{k>i} F_{\beta=k} \right)^{z-1-\psi} \right\} \right) \right]. \tag{12}$$

Again, for the purposes of these expected payoffs, any frequency of empty slots would be incorporated into the total frequency of partners less mutualistic than the focal $\beta=i$ (i.e. into the $F_{\beta=k}$ term in Eq. (12)). In each equation, the term in square brackets gives the probability of failing to form an association. The

first summation in each equation cycles through each partner strategy with which a deal is possible, and is immediately followed by that strategy's frequency (i.e. the probability of interacting with that strategy). The terms in curly brackets reduce the probability of association between two players due to an alternate guild B player being chosen. In Eq. (12), ψ gives each potential additional number of slots that may also be occupied by the same strategy ($\beta=i$), so that when $\psi=0$ the portion of Eq. (12) following the summation over ψ gives the probability of all other slots being occupied by less mutualistic strategies ($\beta=k$), when $\psi=1$ it gives the probability of all other slots but one being occupied by less mutualistic strategies, multiplied by the 50% probability that the focal player is associated with rather than the additional $\beta=i$ player present, and so on.

2.4. Replicator dynamics

The expected payoffs given by Eqs. (11) and (12) can be placed into discrete replicator equations, where for any strategy x of q total strategies in either guild:

$$F_{x(t+1)} = F_{x(t)} \left(\frac{P_{x(t)}}{\sum_{y=1}^q F_{y(t)} P_{y(t)}} \right). \tag{13}$$

We ran simulations using these replicator equations, with $z=2$ members of guild B (i.e. two 'slots') per member of guild A, various initial distributions of demands from 0.01 to 0.99 in each guild, and the assumption that when a successful association occurred surplus equaled one minus the sum of the demands (which also means that associations do not occur when demands sum to an amount greater than one). Discrete versions of the equations were used to allow a more exact match to computer simulations. Simulations were run in Matlab, and strategies were held above a

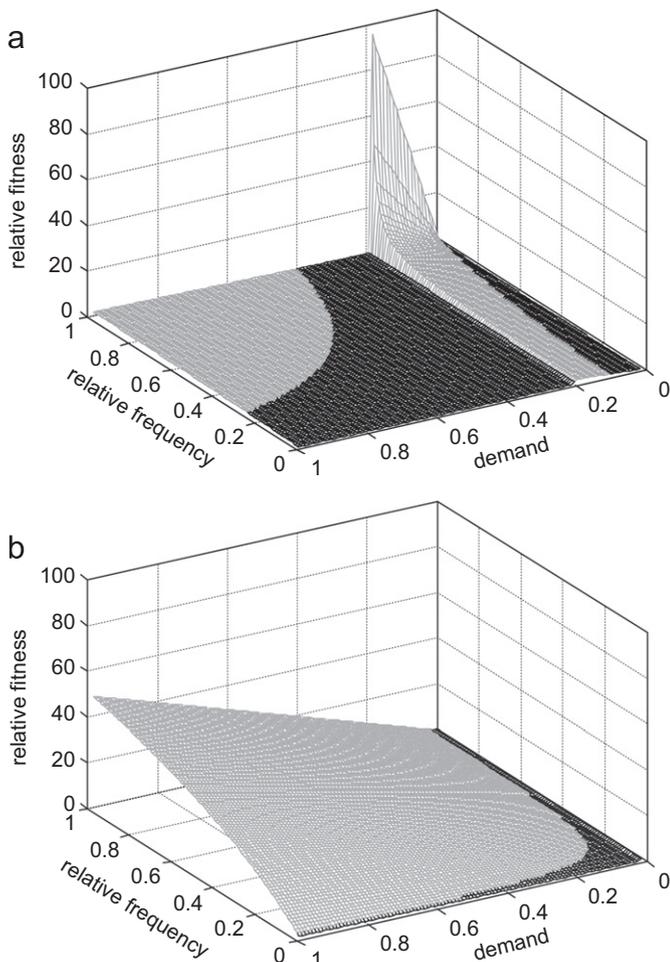


Fig. 1. Expected payoff (fitness) of a focal guild *B* species relative to that of a sole competitor, for two competitor demands (a) 0.2 and (b) 0.01. There are no empty sites, so that relative frequencies of the two competitors sum to one. Guild *A* engages in partner choice and obtains all of the surplus ($s=1$), and is assumed to make all deals offered. Regions where the focal species has a competitive advantage (relative fitness > 1) are shaded gray; black shading indicates relative fitness of 1 or less.

minimum frequency of 10^{-100} (partly to prevent them from going to zero because of rounding error, and partly to keep cycles moving quickly enough to facilitate generating movies to illustrate the dynamics; an extremely low value was chosen to emphasize that only very minimal background frequencies of strategies need be present to generate the results we present). These confirmed the key results drawn from Eqs. (1)–(12). Without any partner choice, the guild receiving the surplus is at a disadvantage (Fig. 2a). If the guild engaging in partner choice also obtains the surplus, highly mutualistic strategies within that guild gain an advantage and persist at high frequency (Fig. 2b). Those mutualistic strategies would be selected against in a guild that engaged in partner choice and gave away the surplus (Fig. 2c), and such a guild would achieve the greatest payoffs. Some strategies persist partly because the dynamics in this system can become extremely slow. Strategies with the minimal demands $a(D_\alpha)=0.01$ in Fig. 2b and $b(D_\beta)=0.01$ in Fig. 2a, for example, have the highest fitness strictly speaking but are so close in fitness to the range of strategies near them that strategy frequencies eventually remain effectively constant. This maintains more demanding strategies in the guild receiving the surplus (as in Fig. 2a), which limits the demand selected for in the opposite guild. Only if the latter guild can also engage in partner choice can it drive its partners to be maximally mutualistic.

When guild *A* both retains all of the surplus and engages in partner choice, highly mutualistic strategies persist in both guilds. The mutualistic strategies in guild *A* present an opportunity for exploiters in guild *B*, but as predicted by Eqs. (9) and (10), those exploiters are effectively lost from the system because they cannot increase from arbitrarily low frequency without some mechanism that generates situations where guild *A* members have no better option than to associate with them (Fig. 2b). Introducing a fixed frequency of empty slots provides exploiters with those situations, and when empty slots are included the result can be patterns of guild *B* strategy frequencies reminiscent of taxon cycles (as explored in Rummel and Roughgarden, 1983, 1985; Brown and Vincent, 1987; Abrams, 1989, 1990; Taper and Case, 1992; Matsuda and Abrams, 1994) or some competition-colonization tradeoff models (Lehman and Tilman, 1997).

The cause of these cyclical dynamics is straightforward, and stems from the analytical results. When one or a group of intermediate strategies is extremely abundant (Fig. 3a), individuals more mutualistic than those almost always succeed in forming associations, and have expected payoffs approaching their demands. Less mutualistic strategies essentially only form associations when they co-occur with an empty slot, so their expected payoffs equal their demand multiplied by the frequency of empty slots. Expected payoffs of the abundant strategies are reduced by the high probability of two individuals of the same strategy co-occurring with the same partner, at which point only one is associated with at random. The strategy slightly more mutualistic than the abundant one(s) therefore has the highest fitness, and so increases in frequency, which leads to the next most mutualistic strategy having highest fitness, and increasing, and so on. This selects for strategies with lower and lower demands, until the expected payoffs of the abundant strategies fall below the payoff that the most exploitative strategy expects purely from co-occurring with empty slots (Fig. 3b). The exploitative strategy's resulting increase in frequency feeds back positively on its expected payoff (due to its increasing chances of co-occurring with itself rather than with a more mutualistic strategy), so that it quickly becomes highly abundant (Fig. 3c), at which point slightly more mutualistic strategies gain an advantage, and the cycle repeats. These patterns occur for various sets of initial guild *A* and *B* strategy frequencies (supplementary material figures S5–S7), and when different sets of simplifying assumptions are used (supplementary material figure S8).

Results in video format associated with this article can be found in the online version at [doi:10.1016/j.jtbi.2010.09.023](https://doi.org/10.1016/j.jtbi.2010.09.023).

3. Discussion

We show that a high degree of mutualism in both guilds is compatible with control of the interaction being highly skewed towards one guild that does what is in its own short-term interests. This occurs because a guild that has control in the sense of being able to appropriate surpluses experiences selective pressure to accept all deals, and a guild that has control in the sense of being able to engage in partner choice is often able to exert selective pressure on its partners to offer increasingly generous deals. That skew of control may also promote diversity by allowing a niche for highly exploitative partners in the guild without control. As in earlier related models (Gale et al., 1995; Toquenaga and Suzuki, 2005; Poulsen, 2007), a guild whose members had the ability to appropriate surpluses but instead gave them away would benefit by having highly mutualistic partners, but within-guild selection would favor members who keep more of the surplus so that the guild would evolve to keep as much of the surplus as possible.

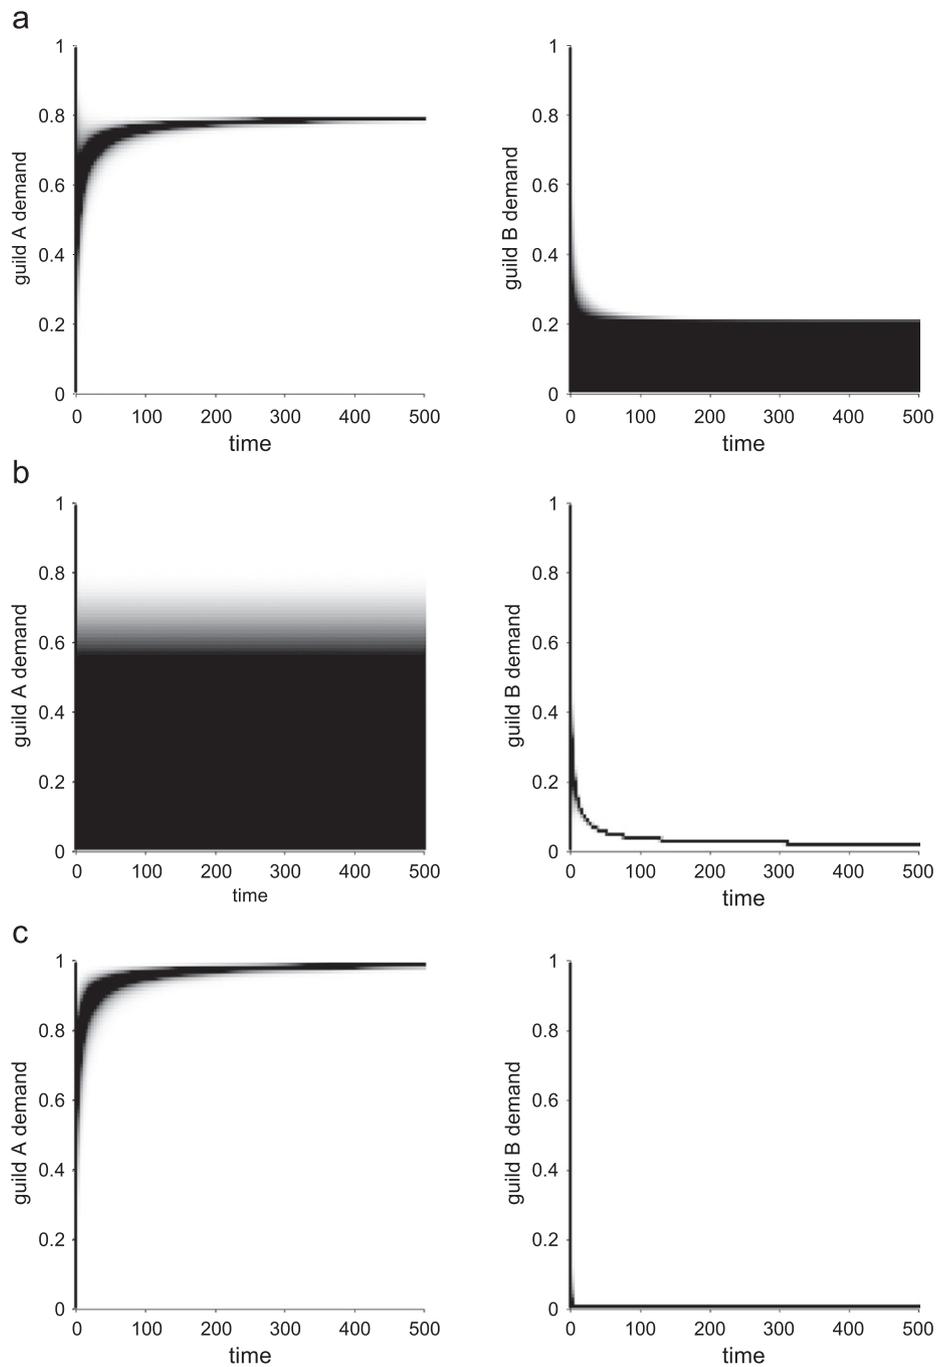


Fig. 2. Strategy frequencies over time in replicator equations with no empty slots under three different scenarios (all strategies have equal initial frequencies in each). Shading gives relative frequency within a timestep, from black=as abundant as the most common strategy to white=absent. Simulations assume that the surplus gained equals one minus the sum of the demands all multiplied by s for the guild A player or $(1-s)$ for the guild B player, that demands must sum to ≤ 1 for a successful association to occur, and that failure to form an association results in zero payoff ($\omega_A = \omega_B = 0$). (a) No choice, surplus to guild B, (b) choice by guild A, surplus to guild A and (c) choice by guild A, surplus to guild B.

Our model assumes one-shot bidding, and no chances for partners that fail to associate because their demands are incompatible to either negotiate or seek alternate partners whose demands may be lower. This represents one simplest case scenario where genetically fixed traits determine how generous versus exploitative a species is based on how demanding it is when forming associations. Our assumption does not preclude the possibility of prolonged associations whose costs and benefits are continually appraised and/or adjusted, merely that the effects of such processes can be integrated over the duration of an interaction to yield an expected payoff to each strategy. For

example, [Johnstone and Bshary \(2002, 2008\)](#) model interactions in which clients terminate interactions with cleaner fish when client benefits (which steadily decrease as fewer parasites remain to be removed) cease to exceed their costs (related to the rate at which a given cleaner takes bites of healthy client tissue). In those models visit duration and payoffs to each partner can be predicted from the start of an interaction. In [Akçay and Roughgarden \(2007\)](#), where negotiation between partners is explicitly modeled, expected average payoffs can again be predicted by initial conditions.

The key aspect of the bargaining process in our model is the potential for incompatible types between which no association

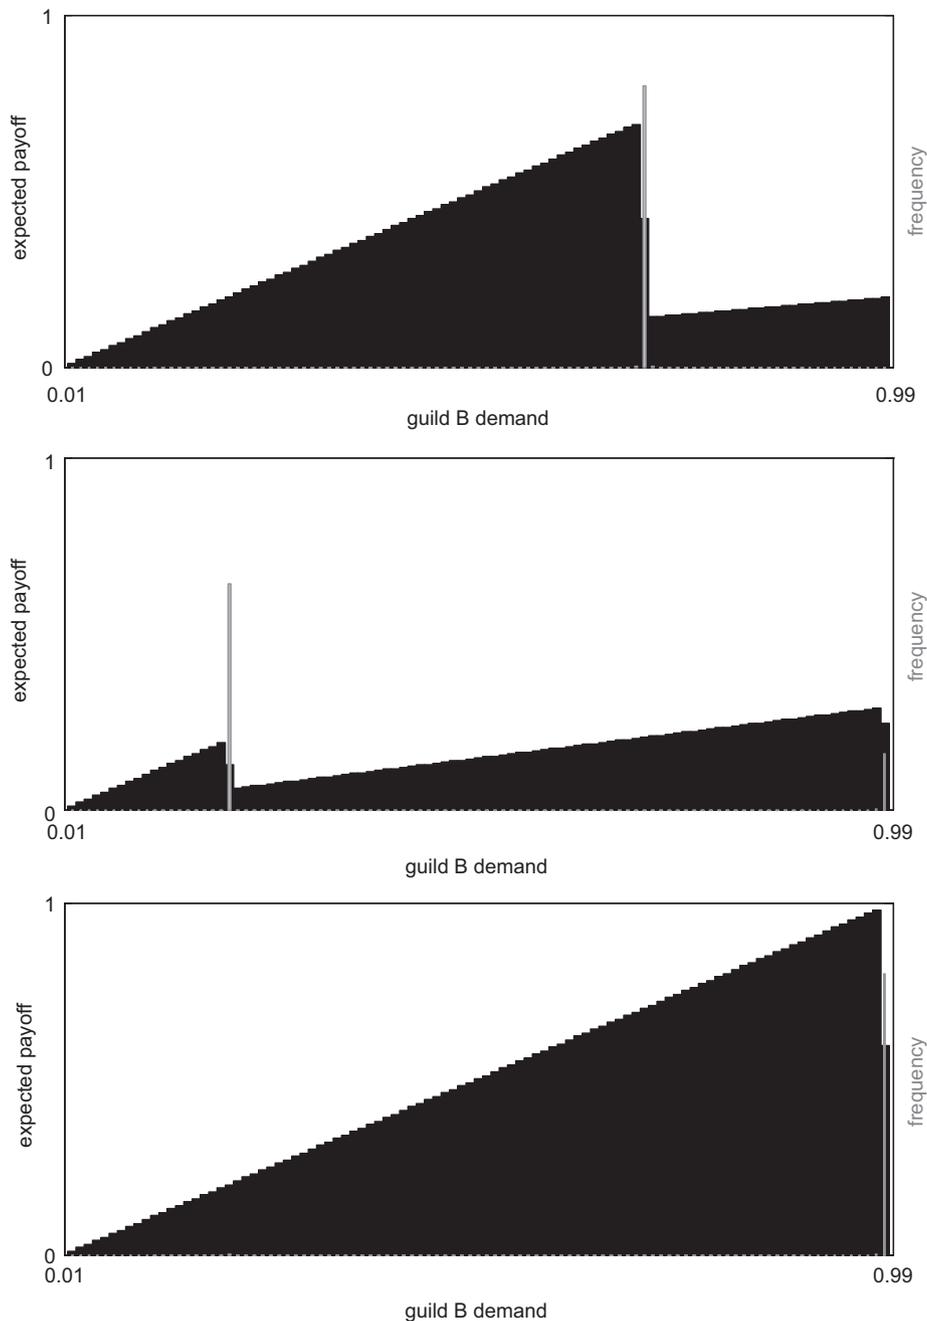


Fig. 3. Cyclical pattern of selection on guild *B* under partner choice, with a fixed frequency of empty slots of 0.2, an even initial distribution of guild *B* demands from 0.01 to 0.99, $s=1$, and all guild *A* players assumed to accept all deals for simplicity. Time: (a) 24820, (b) 31970 and (c) 31980.

will occur. The case that we focus on (where $s=1$ so that guild *A* players receive all surpluses) is consistent with guild *A* players accepting or rejecting offers made by guild *B* players. Because we assume that forming any association is preferable to forming no association, any rejection of an offer by a guild *A* player is disadvantageous and therefore irrational. In this light it is unsurprising that guild *A* evolves towards accepting all offers. However, players in bargaining games may often benefit from having reduced freedom to behave optimally (e.g. see Bergstrom and Lachmann, 2003 and references therein), so it is not necessarily trivial to explore the evolutionary forces acting on fixed irrational strategies. Players effectively behave rationally and accept offers/form associations whenever it is advantageous to do so in most trade-based models of mutualism. This assumption is implicit in models that do not address the division

of benefits between partners, exploring instead when it is selectively advantageous to form any association at all (Schwartz and Hoeksema, 1998; Hoeksema and Schwartz, 2003; de Mazancourt and Schwartz, 2010). Johnstone and Bshary (2002, 2008) assume that cleaner fish never turn away clients, and that clients interact with cleaners for as long as the payoff from doing so is greater than the average payoff they expect once they break off the interaction. In that scenario, if a cleaner bit at a sufficiently high rate, the optimal duration of a visit would be zero, so that sufficiently exploitative types in the guild experiencing choice would fail to form associations. In Noë and Hammerstein (1994), fictional ‘boa constructors’ only turn away less mutualistic ‘shadowbirds’ when they have a sufficiently high probability of subsequently interacting with one that is more generous. Again, when snakes can afford to be choosy, exploitative ‘shadowbirds’

may fail to form associations without an opportunity to negotiate. Other models assume that associations always occur, whether they are profitable or not (e.g. Doebeli and Knowlton, 1998; Ferriere et al., 2002; Bergstrom and Lachmann, 2003; Friesen and Mathias, 2010). In the negotiation process assumed by Akçay and Roughgarden (2007), one partner temporarily halts trade in response to unfavorable changes in the other's behavior regardless of the current benefit being derived, but the cessation is temporary and leads to little lost interaction time overall. The fact that receiving more of the surplus imposes selective pressure on a guild to accept a wider range of offers in our model (see also supplementary figures S4 and S8) illustrates that playing a bargaining role that benefits a species at the level of individual interactions may in and of itself help maintain mutualism in that species. This point has not been well emphasized in the mutualism literature to date.

Our assumption of no negotiation is also consistent with the common framing of the question of generous mutualists' persistence against 'cheaters'. For guilds which make offers and experience choice, strategies are generally envisaged as not being flexible in response to partner pressure. In the models we have discussed so far these inflexible strategies include cleaner fish bite rate (Johnstone and Bshary, 2002, 2008), rhizobial N fixation rate (Friesen and Mathias, 2010), shade provided by 'shadowbirds' (Noë and Hammerstein, 1994), and benefit provision by mycorrhizal fungi (Hoeksema and Kummel, 2003). Negotiation would introduce the possibility that species which are less generous mutualists may be coerced into becoming more generous, so that differences between strategies would have to represent initial rather than final demands.

Expanding our model to explicitly consider subsequent opportunities to form associations when demands are initially incompatible (either due to negotiation or opportunities to interact with alternate potential partners) would be straightforward. The payoffs that would be obtained from such subsequent associations would be multiplied by the probabilities of them occurring for each strategy, and added to Eqs. (11) and (12). Any potential for subsequent associations could reduce the imperative for low demands by choosing guild players (that receive surpluses); this could lower expected payoffs for exploiters in the guild experiencing choice, by reducing the availability of partners willing to interact with exploiters. However, it is reasonable to expect that the value of such subsequent associations might often be reduced, due to time wasted during the failed interaction with the first partner and time required to either find a new partner or engage in a negotiation process (especially if it comes at the expense of the duration of the subsequent successful association). If such reductions occur and are large enough, they could preserve the selective pressure for low demands in the guild obtaining the surplus. Any delay in identifying partner quality (e.g. if partners must be associated with briefly in order to determine their quality) could benefit exploiters in the guild experiencing choice, but payoffs arising from such erroneous associations could be small if the duration is short relative to the duration of a full association.

Separately from our results concerning surpluses, our model yields insights into how partner choice should act as a selective force upon a guild of mutualists. Our main result here is that partner choice may be less effective against highly exploitative types than moderately exploitative types, and may be less effective when highly generous types are abundant. If the vast majority of a guild's members have strategies that lie in some intermediate range, players that are only slightly more exploitative (less mutualistic) than those pay as great a cost in terms of being discriminated against (i.e. not chosen) by partners as the most exploitative players do, but do not reap as high a reward as

the most exploitative players when they do form an association. Thus, selection is strongest against moderately exploitative strategies. Similarly, players that are only slightly more mutualistic than the vast majority of their competitors reap the full rewards of partner choice, but do not pay as high a price to be chosen as the most mutualistic players, so that selection favors moderately mutualistic strategies. These consequences follow naturally from a guild in control choosing the best deals from among those available. They can lead to dynamics reminiscent of taxon cycles seen in other models (see Rummel and Roughgarden, 1983, 1985; Brown and Vincent, 1987; Abrams, 1989, 1990; Taper and Case, 1992; Matsuda and Abrams, 1994), wherein more and more mutualistic strategies are selected for, up to a tipping point at which highly exploitative strategies gain an advantage, and the process repeats. This behavior is also similar to what has been seen in models exploring competition-colonization tradeoffs (Lehman and Tilman, 1997), with mutualism effectively acting as competitive superiority because better mutualists are preferentially chosen by partners (as suggested by Ferriere et al., 2002). Our model highlights that the persistence of highly exploitative types need not be inconsistent with effective discrimination between partners by members of the opposite guild. To the contrary, if choice imposes a uniform deterrent across a range of exploitative strategies, the most highly exploitative of them might be the most able to make sufficient gains to outweigh that deterrent.

Similar to the giving away versus keeping of surpluses, partner choice in these models sets up a situation where a guild could potentially benefit from individuals not doing what is in their own short-term interests. A choosing guild whose members stopped discriminating between partners with demands below a critical threshold could prevent the cyclic dynamic, and over time achieve higher average payoffs. However, selection within such a guild would favor players that ignored such a threshold.

In our simulations, a fixed frequency of empty slots gave exploitative strategies in the guild subjected to choice the low background probability of successful associations necessary for the cyclical dynamics to occur. Empty slots would constitute a relative abundance of partners in the choosing guild, and earlier models incorporating partner choice have highlighted the importance of availability of partners in each guild (Noë and Hammerstein, 1994; Johnstone and Bshary, 2002, 2008). Alternate mechanisms such as spatial clustering, occasional errors by the choosing guild, or invasion at high initial frequency would also be expected to generate the scenarios (where exploiters are chosen regardless of their high demands by a sufficiently high frequency of partners) that would allow such cycles. In our simulations, frequencies of rare strategies decreased with each cycle, so that without some mechanism maintaining minimum strategy frequencies, the cycles would have continually slowed, or stopped if strategies were lost. In real systems, mutation, arbitrary changes in behavior, or immigration from a regional species pool might provide such a mechanism. Since strategies can increase from arbitrarily low frequencies when conditions favor them, the background frequency required is very low (as demonstrated by the 10^{-100} minimum frequency our simulations imposed).

It is interesting to note that in a metapopulation framework, empty slots would be most common when the guild colonizing those slots is highly mutualistic, because it is then receiving low payoffs and would have a low colonization rate. This is exactly when highly exploitative types would be able to take greatest advantage of those empty slots to increase in frequency and generate the cycles seen here. Ironically, in such a scenario exploiters might actually save the system from the 'evolutionary suicide' (as described by Ferriere et al., 2007) that could occur if extremely mutualistic strategies were unable to

maintain a colonization rate that balanced their extinction rate (also analogous to the ‘runaway evolution towards extinction’ (Matsuda and Abrams, 1994) in models of asymmetrical competition that may result in taxon cycles).

Our model is abstract and is intentionally kept simple, but our results are robust to many of our simplifying assumptions (e.g. supplementary material figure S8). The qualitative effects of relaxing many of our simplifying assumptions are straightforward to predict based on how each change would affect the relative prospects for exploiters in the guild experiencing choice to form associations. If reduced chances for highly exploitative types cause the product (payoff per association \times probability of a successful association) to be lower for those exploiters than for highly mutualistic strategies when those mutualists are abundant (i.e. if exploiters do not stand to gain enough relative to highly generous mutualists per successful association), exploiters will be unable to increase in frequency and displace the mutualists. If increased chances cause that product to be higher for slightly more exploitative types than slightly more mutualistic strategies (i.e. if the minimum amount of payoff that a mutualist must tradeoff to be preferentially chosen over an exploiter is too great), the exploiters will not be displaced and partner choice will not be capable of selecting for greater and greater mutualism. More ‘slots’ per chooser, which would allow choosers to sample a greater number of potential partners, should reduce the chances for exploiters in the chosen guild to form associations, although if choosers associated with more than one of those potential partners this could help exploiters. We assume that choice imposes an equal penalty on any rare strategy that is more exploitative than other strategies present; sanctions that punish partners proportional to their greed could reduce the advantage that highly exploitative types experience over moderately exploitative types. Similarly, added rewards for highly generous partners could favor them over more moderately generous partners. However, if imposing those sanctions or providing those rewards also reduces the benefit gained from association with highly exploitative or highly generous types, it would be disadvantageous for choosing guild players to impose them (since we already assume exploiters are associated with only when no better option exists). Our assumption that the more mutualistic of two potential partners is chosen no matter how slight the difference between the two is an extreme one, but this assumption has the effect of facilitating the potential for partner choice to select for mutualistic partners, which is the conventional view of partner choice’s effect. The assumption does not play into our more novel main result that choice can be less effective against extremely exploitative types than against moderately exploitative ones. Further, all that is required for partner choice to select for more and more generous types is for the cost of whatever difference in generosity is required to be preferentially chosen by partners to be sufficiently small relative to the benefit of being preferentially chosen. The potential of any of these added considerations to alter selective pressures in our model depends on their quantitative effects on the expected payoffs to competing strategies in each guild. In any case, receiving surpluses should still lead to lower demands than would be the case otherwise, and highly exploitative types should still experience an advantage over moderately exploitative types that is most pronounced when their competitors are predominantly highly mutualistic.

The broad outlines of the framework used here are potentially relevant to several mutualistic systems. As stated in the introduction, the dependence of both payoffs on guild *B*’s demand is consistent with a scenario where benefits to guild *B* and the costs to guild *A* constitute a consumer-resource interaction. Thus guild *B*’s demand might represent the carbon (energy) gain

demand by a mycorrhizal or rhizobial symbiont for a unit of nutrient, an ant for a unit of protection, or a pollinator or cleaner fish for its service. There are reasons why the conflict of interests we assume might not always be present. A cleaner fish, for example, obtains energy in the course of providing its service (since it is eating the parasites it is removing), so that there may be situations where increasing benefits to the cleaner also mean increasing benefits to the client. However, cleaners may also consume tissue and mucus from the client (Bshary, 2001), which is more likely to benefit the cleaners at the client’s expense and introduce the kind of conflict of interests we consider here. In our model, interests would be aligned when a change in the terms of an interaction benefit both parties. In this case, the interaction terms that are worse for both parties should not be seen unless some type of constraint prevents interacting pairs from adopting the terms that are superior for both. Because any alignment of interests is unlikely when comparing deals at the extreme limits of profitability for each side, it is much more likely to limit the (more conventionally accepted) potential of partner choice to select for mutualism than the more novel advantage to extremely exploitative types that we suggest. Our framework also does not consider potential longer-term feedbacks. A highly mutualistic strategy, for example, might benefit down the road from increasing the local abundance of partners and/or their capacity to deliver benefits. In real systems, such mechanisms could operate alongside and interact with those we explore here, but they are outside the scope of our current model.

Viewing control of mutualisms explicitly in terms of each partner’s ability to appropriate surpluses could allow alternate perspectives on some interactions, and may facilitate comparisons among interaction types. For example, Johnstone and Bshary (2002, 2008) model a cleaner fish mutualism in which they argue that clients’ abilities to determine the durations of interactions counterbalances cleaners’ one-sided ability to exploit the interaction by taking bites of healthy client tissue. However, if a range of biting rates yields interactions that remain profitable for both parties, it could be argued that extremely low rates of biting would constitute exploitation of cleaners by clients just as much as extremely high rates constitute exploitation of clients by cleaners. Considering the system in terms of minimum demands and surpluses explains the intermediate bite rates that their model often predicts. If clients are willing to interact with a given cleaner strategy for any nonzero duration, cleaners with very slightly higher bite rates would be interacted with for some shorter duration. This means that the client fish obtains higher payoff than the absolute minimum it demands. Similarly, a cleaner that interacts with a client for a given duration would have been (and was) willing to interact for a shorter duration, so that the cleaner also obtains a higher payoff than the minimum it demands. Thus, the above scenario is consistent with $0 < s < 1$, which predicts that neither guild evolves to be maximally mutualistic or maximally exploitative. Alternately, cleaners in the model could be seen as simply proposing a deal and clients responding, since for any given community context (cleaner and client abundances and strategy frequencies), visit duration and thus payoff to each player is a function of bite rate. From that perspective, clients obtain the entire surplus and cleaners should be free to demand the maximum payoff possible under that constraint (i.e. choosing the bite rate that is most profitable after accounting for its effect on visit duration).

Our work stresses that consideration should be given to potential downsides as well as potential benefits of control over mutualistic interactions (and the ability to maximize short-term payoffs in them). Partner choice in particular, especially when acting in concert with other mechanisms, may have richer implications for mutualism than widely appreciated. Because exploiters in our model rely on

situations where a partner has no other option but to associate with them, they would be favored by factors such as high initial frequency, empty patches, or spatial clumping, which all favor mutualism in many other models (e.g. reviewed by Lion and van Baalen, 2008). This represents a role-reversal for the two strategies, and suggests that we should not always assume that mutualism is the trait in need of further explanation.

Acknowledgements

A. Golubski was supported during this work by grants to K.L. Gross and G.G. Mittelbach from the AW Mellon Foundation and by NSF postdoctoral fellowship OISE-0754419. C. Klausmeier was supported during this work by the NSF and James S McDonnell Foundation. Thanks to Erol Akçay, Jay Biernaske, Jessica Forrest, Andy Gardener, and anonymous reviewers for feedback on earlier drafts of the manuscript. This is KBS contribution number 1574.

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.jtbi.2010.09.023.

References

- Abrams, P.A., 1989. The importance of intraspecific frequency-dependent selection in modelling competitive coevolution. *Evol. Ecol.* 3, 215–220.
- Abrams, P.A., 1990. Ecological versus evolutionary consequences of competition. *Oikos* 57, 147–151.
- Akçay, E., Roughgarden, J., 2007. Negotiation of mutualism: rhizobia and legumes. *Proc. R. Soc. London B* 274, 25–32.
- Bergstrom, C.T., Lachmann, M., 2002. The red king effect: evolutionary rates and the division of surpluses in mutualisms. In: Hammerstein, P. (Ed.), *Genetic and Cultural Evolution of Cooperation*. MIT Press, Cambridge, MA, pp. 223–238.
- Bergstrom, C.T., Lachmann, M., 2003. The red king effect: when the slowest runner wins the coevolutionary race. *Proc. Natl. Acad. Sci. USA* 100, 593–598.
- Bever, J.D., Richardson, S.C., Lawrence, B.M., Holmes, J., Watson, M., 2009. Preferential allocation to beneficial symbiont with spatial structure maintains mycorrhizal mutualism. *Ecol. Lett.* 12, 13–21.
- Brown, J.S., Vincent, T.L., 1987. Coevolution as an evolutionary game. *Evolution* 41, 66–79.
- Bshary, R., 2001. The cleaner fish market. In: Noë, R., van Hooff, J.A.R.A.M., Hammerstein, P. (Eds.), *Economics in Nature: Social Dilemmas, Mate Choice and Biological Markets*. Cambridge University Press, Cambridge, pp. 146–172.
- Bshary, R., Grutter, A.S., 2002. Experimental evidence that partner choice is a driving force in the payoff distribution among cooperators or mutualists: the cleaner fish case. *Ecol. Lett.* 5, 130–136.
- Bshary, R., Schäffer, D., 2002. Choosy reef fish select cleaner fish that provide high-quality service. *Animal Behav.* 63, 557–564.
- Bull, J.J., Rice, W.R., 1991. Distinguishing mechanisms for the evolution of cooperation. *J. Theor. Biol.* 149, 63–74.
- Denison, R.F., 2000. Legume sanctions and the evolution of symbiotic cooperation by rhizobia. *Am. Nat.* 156, 567–576.
- Doebeli, M., Knowlton, N., 1998. The evolution of interspecific mutualisms. *Proc. Natl. Acad. Sci. USA* 95, 8676–8680.
- Edwards, D.P., Hassall, M., Sutherland, W.J., Yu, D.W., 2006. Selection for protection in ant-plant mutualism: host sanctions, host modularity, and the principal-agent game. *Proc. R. Soc. B* 273, 595–602.
- Ferriere, R., Bronstein, J.L., Rinaldi, S., Law, R., Gauduchon, M., 2002. Cheating and the evolutionary stability of mutualisms. *Proc. R. Soc. London B* 269, 773–780.
- Ferriere, R., Gauduchon, M., Bronstein, J.L., 2007. Evolution and persistence of obligate mutualists and exploiters: competition for partners and evolutionary immunization. *Ecol. Lett.* 10, 115–126.
- Fischer, S., Güth, W., Müller, W., Stiehler, A., 2006. From ultimatum to Nash bargaining: theory and experimental evidence. *Exp. Econ.* 9, 17–33.
- Foster, K.R., Wenseleers, T., 2006. A general model for the evolution of mutualisms. *J. Evol. Biol.* 19, 1283–1293.
- Friesen, M.L., Mathias, A., 2010. Mixed infections may promote diversification of mutualistic symbiotes: why are there ineffective rhizobia? *J. Evol. Biol.* 23, 323–334.
- Gale, K., Binmore, K.G., Samuelson, L., 1995. Learning to be imperfect: the ultimatum game. *Game Econ. Behav.* 8, 56–90.
- Güth, W., Schmittberger, R., Schwarze, B., 1982. An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Organ.* 3, 367–388.
- Herre, E.A., Knowlton, N., Mueller, U.G., Rehner, S.A., 1999. The evolution of mutualisms: exploring the paths between conflict and cooperation. *Trends Ecol. Evol.* 14, 49–53.
- Hoeksema, J.D., Kummel, M., 2003. Ecological persistence of the plant-mycorrhizal mutualism: a hypothesis from species coexistence theory. *Am. Nat.* 162, S40–S50.
- Hoeksema, J.D., Schwartz, M.W., 2003. Expanding comparative-advantage biological market models: contingency of mutualism on partners' resource requirements and acquisition trade-offs. *Proc. R. Soc. London B* 270, 913–919.
- Johnstone, R.A., Bshary, R., 2002. From parasitism to mutualism: partner control in asymmetric interactions. *Ecol. Lett.* 5, 634–639.
- Johnstone, R.A., Bshary, R., 2007. Indirect reciprocity in asymmetric interactions: when apparent altruism facilitates profitable exploitation. *Proc. R. Soc. London B* 274, 3175–3181.
- Johnstone, R.A., Bshary, R., 2008. Mutualism, market effects, and partner control. *J. Evol. Biol.* 21, 879–888.
- Kiers, E.T., Rosseau, R.A., West, S.A., Denison, R.F., 2003. Host sanctions and the legume-rhizobium mutualism. *Nature* 425, 78–81.
- Lehman, C.L., Tilman, D., 1997. Competition in spatial habitats. In: Tilman, D., Kareiva, P. (Eds.), *Spatial Ecology: the Role of Space in Population Dynamics and Interspecific Interaction*. Princeton University Press, Princeton, pp. 185–203.
- Lion, S., van Baalen, M., 2008. Self-structuring in spatial evolutionary ecology. *Ecol. Lett.* 11, 277–295.
- Matsuda, H., Abrams, P.A., 1994. Runaway evolution to self-extinction under asymmetrical competition. *Evolution* 48, 1764–1772.
- de Mazancourt, C., Schwartz, M.W., 2010. A resource ratio theory of cooperation. *Ecol. Lett.* 13, 349–359.
- Nash, J.F., 1950. The bargaining problem. *Econometrica* 18, 155–162.
- Noë, R., 2001. Biological markets: partner choice as the driving force behind the evolution of mutualisms. In: Noë, R., van Hooff, J.A.R.A.M., Hammerstein, P. (Eds.), *Economics in Nature: Social Dilemmas, Mate Choice and Biological Markets*. Cambridge University Press, Cambridge, pp. 93–118.
- Noë, R., Hammerstein, P., 1994. Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism, and mating. *Behav. Ecol. Sociobiol.* 35, 1–11.
- Noë, R., Hammerstein, P., 1995. Biological markets. *Trends Ecol. Evol.* 10, 336–339.
- Noë, R., van Schaik, C.P., van Hooff, J.A.R.A.M., 1991. The market effect: an explanation for pay-off asymmetries among collaborating animals. *Ethology* 87, 97–118.
- Poulsen, A.U., 2007. Information and endogenous first mover advantages in the ultimatum game: an evolutionary approach. *J. Econ. Behav. Organ.* 64, 129–143.
- Rummel, J.D., Roughgarden, J., 1983. Some differences between invasion-structured and coevolution-structured competitive communities: a preliminary theoretical analysis. *Oikos* 41, 477–486.
- Rummel, J.D., Roughgarden, J., 1985. A theory of faunal buildup for competition communities. *Evolution* 39, 1009–1033.
- Sachs, J.L., Mueller, U.G., Wilcox, T.P., Bull, J.J., 2004. The evolution of cooperation. *Q. Rev. Biol.* 79, 135–160.
- Schwartz, M.W., Hoeksema, J.D., 1998. Specialization and resource trade: biological markets as a model of mutualisms. *Ecology* 79, 1029–1038.
- Simms, E.L., Taylor, D.L., 2002. Partner choice in nitrogen-fixation mutualisms of legumes and rhizobia. *Integ. Comp. Biol.* 42, 369–380.
- Simms, E.L., Taylor, D.L., Povich, J., Shefferson, R.P., Sachs, J.L., Urbina, M., Tausczik, Y., 2006. An empirical test of partner choice mechanisms in a wild legume-rhizobium interaction. *Proc. R. Soc. London B* 273, 77–81.
- Taper, M.L., Case, T.J., 1992. Models of character displacement and the theoretical robustness of taxon cycles. *Evolution* 46, 317–333.
- Toquenaga, Y., Suzuki, Y., 2005. Evolution of generosity in the demand game. *Pop. Ecol.* 47, 31–39.
- West, S.A., Kiers, E.T., Simms, E.L., Denison, R.F., 2002. Sanctions and mutualism stability: why do rhizobia fix nitrogen? *Proc. R. Soc. London B* 269, 685–694.
- West, S.A., Diggle, S.P., Buckling, A., Gardner, A., Griffins, A.S., 2007. The social lives of microbes. *Annu. Rev. Ecol. Syst.* 38, 53–77.