Walrasian equilibrium behavior in nature

Ted Loch-Temzelides[a,1]

[a]Department of Economics, Rice University, Houston, TX 77005

Edited by Stephen Polasky, University of Minnesota, St. Paul, MN, and approved April 16, 2021 (received for review October 12, 2020)

The interaction between land plants and mycorrhizal fungi (MF) forms perhaps the world's most prevalent biological market. Most plants participate in such markets, in which MF collect nutrients from the soil and trade them with host plants in exchange for carbon. In a recent study, M. D. Whiteside et al. [Curr. Biol. 29, 2043–2050.e8 (2019)] conducted experiments that allowed them to quantify the behavior of arbuscular MF when trading phosphorus with their host roots. Their experimental techniques enabled the researchers to infer the quantities traded under multiple scenarios involving different amounts of phosphorus resources initially held by different MF patches. We use these observations to confirm a revealed preference hypothesis, which characterizes behavior in Walrasian equilibrium, a centerpiece of general economic equilibrium theory.

Walrasian behavior | biological markets | revealed preference

Several biological interactions concern the direct or indirect exchange of valuable resources. When such exchanges involve human participants, they constitute the main focus of economics. Similar interactions that involve at least two distinct classes of nonhuman participants choosing their trading partners and their individual trading patterns are referred to as biological markets (1–5). While the study of biological markets by biologists often uses concepts borrowed from economics, such as supply and demand, prices, etc., biologists often tend to emphasize the game-theoretic aspects of such interactions (6–8). This seems appropriate since their analysis tends to focus on concepts like reciprocity, cooperation, cheating, exploitation, etc., which are game-theoretic in nature and do not have a direct analog in general economic equilibrium theory (GET). Leading examples in the theory of biological markets can be found in refs. 9–11, while refs. 12–17 are notable laboratory studies of biological markets. An application of contract theory to biological markets can be found in ref. 18. Named after Leon Walras, *Walrasian GET* forms the foundation of the economic study of markets. In the context of a pure exchange of resources, a model economy in GET consists of a set of traders, each characterized by their preferences over consuming different combinations (baskets) of resources (goods), together with a description of the quantities of each good that each trader happens to be endowed with prior to the exchange taking place. Given the exchange rates (prices) among different goods, the traders decide how much of each good to buy, sell, or keep for their own consumption in order to maximize their well-being, subject to their affordability (budget) constraints. The latter require that, for each trader, the value of goods purchased in the market—i.e., their total expenditure—does not exceed the value of the goods they possess prior to trading—i.e., their total initial wealth.

GET is a powerful tool and has a variety of measurable and refutable implications. It prescribes outcomes, known as *Walrasian equilibria* (WE). A formal definition of WE is given in *Materials and Methods*. Intuitively, it consists of a set of prices and a set of final consumption baskets of different goods for each market participant that satisfy two conditions. First, taking prices as given, each trader chooses the basket that maximizes her payoff (utility) among those that are affordable to her. Second, the resulting final allocation of goods among all traders is feasible—i.e., the total amount allocated (demand) equals the total amount existing in the economy (supply). Two classic treatments of GET are Debreu (19) and Arrow and Hahn (20). Much like Nash equilibrium in game theory, WE is the central solution concept in economics. Unlike Nash equilibrium, which deals with strategic interactions among participants in specific institutional contexts (often referred to as noncooperative games), WE describes trading outcomes regardless of the underlying institutional context, provided that markets are "frictionless." The absence of market frictions is defined by the requirements of no information asymmetries about the quality of the exchanged resources, small or no transaction costs, and, importantly, price-taking behavior. The latter implies that traders act as if their own behavior had no effect on the terms of trade. Technically, price-taking behavior requires that traders choose the quantities they trade in the market in order to maximize their own payoff, while treating prices as parameters. For each traded commodity, resource scarcity implies that the total quantity demanded by buyers does not exceed the total quantity supplied by sellers. These feasibility constraints assist in determining prices. The only requirements imposed by WE are that, given prices, the resulting trades are affordable to each trader, and the final allocation of resources is feasible. In other words, GET postulates that participants in frictionless markets act "as if" they followed Walrasian behavior.

To an economist, despite institutional, strategic, and other considerations, all frictionless markets are Walrasian. This implies that, provided that traders act as if they had no market power, and there are small or no other frictions, the final allocation of resources across market participants is the one prescribed by WE. Importantly, this "Walrasian hypothesis" is expected to hold, even if actual trade takes place in ways that do not resemble Walrasian markets (21). A fundamental question that, to our

Significance

Mutualisms are commonly observed ecological interactions, often involving the exchange of resources across species. Such exchanges can be thought of as biological markets. Biologists modeling these markets often employ an informal mix of economics and game-theoretic concepts. A fundamental question is whether exchange in biological markets is consistent with general economic equilibrium theory (GET), the main paradigm used to study exchange in economics. This paper uses data from biological experiments to demonstrate that the trading behavior of mycorrhizal fungi is consistent with the predictions of GET. The large volume of knowledge in GET might result in new insights about biological exchange. In turn, experimental findings in biology can lead to a new field of application for GET.

Author contributions: T.L.-T. performed research and wrote the paper.

The author declares no competing interest.

This article is a PNAS Direct Submission.

[1]Email: tedt@rice.edu.

Published June 28, 2021.

ECONOMIC SCIENCES

ECOLOGY

knowledge, has not been formally addressed is whether biological markets studied from the point of view of GET give rise to behavior that is consistent with WE outcomes. While not easy to answer, this hypothesis can be investigated rigorously. In many markets, including biological ones, it is hard or impossible to directly observe the traders' payoffs over different combinations of goods. However, in biological experiments, we might be able to observe prices as well as the magnitudes of resulting commodities consumed by the traders at different prices. The celebrated Afriat theorem asserts that if in a given market, observations on quantities and prices satisfy a set of inequalities, known as the Generalized Axiom of Revealed Preference (GARP), then they are consistent with WE outcomes.

The interaction between land plants and mycorrhizal fungi (MF) forms perhaps the world's most prevalent biological market. Most plants participate in such markets, in which MF collect nutrients from the soil and trade them with host plants in exchange for carbon (C) (22). In this paper, we use data from biological experiments conducted by Whiteside et al. (23). Although their study was designed with different questions in mind, in the course of their experiment, they carefully document the prices and the corresponding combinations of C and phosphorus (P) obtained by MF after trading with their host roots. Using Afriat's theorem, we demonstrate that the trading behavior of MF in their experiment is consistent with WE.

The classic literature on the game-theoretic foundations of WE relies on lack of market power. However, there are evolutionary arguments in support of WE in finite economies. Alchian (24) argued that evolution would favor relative, rather than absolute, performance. Schaffer (25) demonstrated that the Walrasian outcome is the only evolutionary stable strategy in a static Cournot duopoly model. Vega-Redondo (26) reached a similar conclusion using stochastic evolutionary dynamics in a repeated Cournot oligopoly model. Kim and Wong (27) showed that a Walrasian outcome is the only stochastically stable state in a Shapley–Shubik market game with Leontief preferences (see also ref. 28). Although they accommodate a small number of agents, these studies have some common features that reduce the importance of market power. First, agents are assumed to be myopic. Second, their actions are subject to "mutations"—i.e., vanishingly small levels of noise. Third, the payoffs emphasize relative, as opposed to absolute, performance. These conditions seem rather natural in a biological context, especially one involving fungi.

A couple of remarks are in order. First, WE does not rely on the existence of posted prices, auctioneers, or, indeed, conscious optimization by the traders. It is silent about the negotiating protocol through which prices are reached and about how traders might interact away from equilibrium prices. It only gives a prediction about what the resulting terms of trade and final resource allocation will look like if traders trade in a WE [Hurwicz et al. (29) discuss alternative decentralized allocation mechanisms]. Second, WE outcomes have some desirable efficiency properties. Although every trader is assumed to act in a purely selfish fashion, the resulting allocation of resources under a WE is efficient in the sense that no trader can become better off by a further reallocation of resources, without having to make another trader worse off. This is sometimes referred to as Pareto efficiency; see ref. 19 for a classic axiomatic treatment. Put differently, no gains from trade remain unexploited, and no resources are wasted. Hence, establishing WE behavior provides an indirect proof that, while individually selfish, market participants in the biological market under study reach efficient outcomes for their respective ecosystem taken as a whole. Thus, WE can provide a useful efficiency benchmark, even when game-theoretic models are used to model the underlying exchange.

The paper proceeds as follows. *Biological Markets—An Example* describes the biological market in the Whiteside et al. (23) experiment. *WE Behavior by MF* uses the experimental data to demonstrate compliance with WE behavior. A brief conclusion follows. More technical derivations can be found in *Materials and Methods*.

Biological Markets—An Example

Our analysis is based on data from experiments performed by Whiteside et al. (23). They quantify P-trading behavior of arbuscular MF using quantum-dot tracking techniques. Their experiments consider fungi that are exposed to "rich" and "poor" resource patches in terms of the amount of P available. Both the MF and the hosts interact with multiple partners and can discriminate among them based on the type and amount of resources they need. Whiteside et al. use fluorescent nanoparticles to determine the quantity of P transferred to plant hosts in exchange for C. Remaining P after trade can be used or stored. Stored nutrients can be used or traded by the MF network in the future. Their techniques allow them to infer the P quantities traded under multiple scenarios involving different (unequal) amounts of P resources initially held by the MF. They focused on identifying the effects of the different levels of inequality on trade. They found that MF responded to high resource variation by increasing the total amount of P traded with host roots, decreasing their own consumption (storage), and moving resources within the network, mostly from rich to poor patches, in order to gain better returns from increased demand. Although not designed for this purpose, this experimental design proved essential for our analysis concerning whether the trading behavior of the MF is consistent with WE.

In their experiments, Whiteside et al. kept the total amount of P constant, but varied the total ratio across two fungus compartments. They considered three cases for this ratio: 90:10 ("high inequality"), 70:30 ("medium inequality"), and 50:50 ("no inequality"). They proceeded to measure first the amount of P transferred between the fungal compartments and, subsequently, the amount of P traded by each fungus compartment in exchange for C from their respective plant root. The latter exchange is the focus of our study. While their techniques allow them to quantify fungal transfers, they did not measure C transfers to the MF directly. Instead, they considered fungal biomass as a proxy for the amounts of C transferred by the host roots. In what follows, we will use the same quantity (milligrams of fungal biomass) to measure implied C consumption, assuming that the MF had no other access to C resources. Schematically, our analysis is based on price and quantity data related to the markets shown in Fig. 1. For the purposes of mapping their setup to the GET model, we will consider two separate markets. In market 1, rich MF (high initial endowment of P) exchange P for C with their host root at the observed exchange rate. In market 2, poor MF (low initial endowment of P) trade P for C under the observed exchange rate. We do not model the intranetwork transfers of P between rich and poor MF compartments and will instead focus on the trading in markets 1 and 2 after such transfers. Our focus will be to investigate whether the observed prices and quantities in each of the two markets are consistent with Walrasian behavior by the MF.

Whiteside et al. (23) quantified the P transfer to the hosts' root segments, P retained by the hyphal network, and the implied exchange rates (prices) measured in C (biomass) obtained by the MF per 1 nmol of P. Furthermore, their experiment guaranteed that the economic assumption of nonsatiation holds, as the trading parties were prevented from being overly saturated. Their experimental treatment is uniquely suitable for our analysis, as they consider three different cases regarding the distribution (inequality) of P allocated in different MF patches.

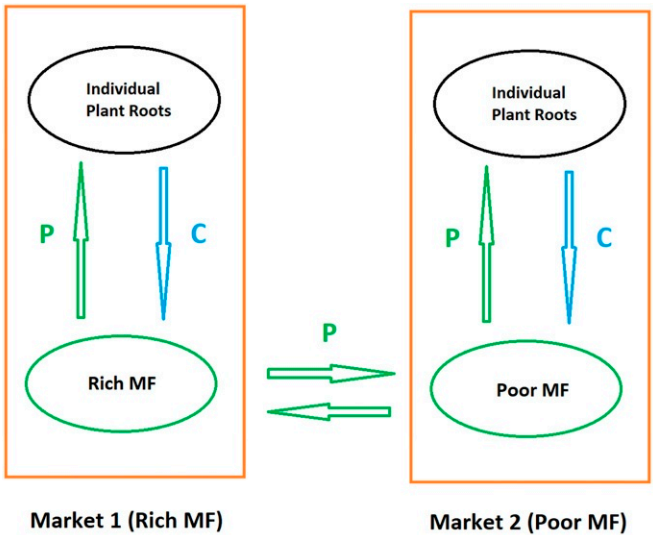

Fig. 1. The two markets—P exchanged for C.

Different prices were observed per inequality level and per resource compartment. These are defined by the ratio of fungal hyphal biomass in milligrams (C received) divided by total amount of P transferred to the respective root. The price data allow us to infer the amounts of C received by the MF in this exchange (*Materials and Methods*). While the original experimental design addressed different questions, these observations are essential for our purposes, as they allow us to test the *Revealed Preference Hypothesis*. Developed in Samuelson (30), Revealed Preference is a central hypothesis in economics. It postulates that observed individual choices must satisfy certain consistency criteria, preventing logical contradictions when ranking different consumption baskets. This hypothesis is a necessary and sufficient condition for WE behavior.

WE Behavior by MF

Viewed through the lens of GET, the experimental data are naturally divided into two markets, as rich and poor MF have access to different endowments of P, and they trade with their host plants under different prices. We will concentrate on whether the resulting trades between MF and the host roots can be rationalized as part of a WE. The experimental design implied that each market was run under three different setups in terms of inequality: none, medium, and high. This resulted in the set of price and demand data in Table 1 (market 1) and Table 2 (market 2) (*Materials and Methods*). For example, price vector $p_1 = (1, 12.80710378)$ indicates a price of C of 12.80710378 per nmol of P (the price of P is normalized to one). Under price p_1, the resulting consumption basket for the MF was $x_1 = (0.054608356, 0.004879507)$. The first number corresponds to the amount of P purchased by the MF (retained quantities count as a purchase in GET, as they imply a marginal cost equal to the corresponding price), while the second number corresponds to the amount of C purchased from the plant in exchange for the P sold by the MF. For the rich MF, we have:

Are the terms of trade and the corresponding baskets chosen by the MF consistent with WE? A powerful result in GET, known as Afriat Theorem, characterizes WE behavior. It applies in circumstances like ours, where a finite set of data observations on prices and corresponding baskets are obtained from actual trading. In other words, Afriat's Theorem offers a set of testable implications that are necessary and sufficient for WE behavior. In order to verify the conditions of the theorem, we will proceed in three steps. References on revealed preference and Afriat's Theorem in GET include Brown and Matzkin (31) and Fostel et al. (32).

If basket x_i is chosen by a trader under price vector p_i, and a different basket x_j is chosen under price p_j, and basket x_j was affordable under price p_i—i.e., $p_i x_j < p_i x_i$—then we say that x_i *is Directly Revealed Preferred to* x_j. We denote this by $x_i \succeq_D x_j$. The *Weak Axiom of Revealed Preference* (WARP) requires that the relationship $\succeq_D$ is asymmetric. In other words, if x_i is directly revealed preferred to x_j, then x_j cannot be directly revealed preferred to x_i.

$$\textbf{WARP}: x_i \succeq_D x_j \Rightarrow \text{not } x_j \succeq_D x_i. \qquad [1]$$

Put simply, if under prices p_i you buy bundle x_i when you could afford bundle x_j, then you reveal that you prefer x_i to x_j. Thus, when under p_j you buy x_j, it must be that x_i is not affordable.

Fig. 2 gives an example for the case of two goods where WARP is violated. To see this, suppose that bundle x_1 is chosen under price p_1, while bundle x_2 is chosen under price p_2. Since x_2 is affordable under p_1, we have that $x_1 \succeq_D x_2$. In addition, since x_1 is affordable under p_2, we have that $x_2 \succeq_D x_1$, contradicting WARP. As we will discuss later, WARP is a necessary condition for WE behavior. Thus, the choices in this example are not consistent with WE.

WARP can be thought of as a minimum criterion for consistency in a trader's choices. It provides an implication that is readily testable given our data and that, to our knowledge, has never before been tested in trade involving nonhuman participants. Yet, WARP does not guarantee behavior consistent with WE. Next, we define $x_0 \succeq x_n$ to mean that there is a chain of x_m such that $x_m \succeq_D x_{m+1}$ for $m = 0, \ldots, n-1$, and we let $\succ$ denote the asymmetric part of the binary relation $\succeq$. The *Strong Axiom of Revealed Preference* (SARP) requires that if $x_i \succeq x_j$, then it cannot be that $x_j \succeq x_i$. In other words, SARP requires a transitivity property assuring that no cyclical contradictions appear in the observed data.

$$\textbf{SARP}: \textit{There is no chain such that } x_0 \succeq_D \ldots \succeq_D x_m \succeq_D x_m + 1 \succeq_D \ldots \succeq_D x_0. \qquad [2]$$

For example, if $x_1 \succeq_D x_2$ and $x_2 \succeq_D x_3$, SARP requires that it cannot be that $x_3 \succeq_D x_1$. If, in addition, traders are not satiated—i.e., they have not reached levels of consumption that are so high that additional consumption would be detrimental, and we assume that there is a single level of consumption that maximizes payoff at each price—SARP is equivalent to the following variation, known as *GARP*:

$$\textbf{GARP}: \textit{If } x \succeq y, \textit{then it cannot be that } y \succ x. \qquad [3]$$

In other words, if x is directly or indirectly revealed preferred to y, then it cannot be that y is strictly directly preferred to x. The nonsatiation assumption is verified to hold in the current experimental setup. This guarantees, among other things, that prices remain strictly positive as additional consumption remains desirable. We are now ready to state Afriat's theorem, the main tool for our analysis. The part of the theorem that is relevant for our purposes can be stated as follows:

Table 1. Market 1 Data—Rich MF

Rich	p_i (C/nmol of P)	P consumed	C consumed
Basket x_1	12.80710378	0.054608356	0.004879507
Basket x_2	6.958750971	0.040385412	0.004015199
Basket x_3	5.155169512	0.010051909	0.004036498

ECONOMIC SCIENCES

ECOLOGY

Similarly, for the poor MF we have:

Table 2. Market 2 Data—Poor MF

Poor	p_i (C/nmol of P)	P consumed	C consumed
Basket x_4	14.2962891	0.056380193	0.004903627
Basket x_5	17.46152694	0.033355445	0.005447996
Basket x_6	24.77908622	0.006889843	0.005872643

Theorem (Afriat). *The following two conditions are equivalent*: 1) *Observations* $(p_1, x_1), \ldots, (p_n, x_n)$ *satisfy GARP*; *and* (2) *the data are consistent with WE behavior for some nonsatiated preferences.*

In summary, GARP requires that the different price–quantity combinations observed satisfy a "consistency" condition similar to a transitivity property. Afriat's theorem asserts that this condition is necessary and sufficient for WE behavior by the respective trader.

As an example, consider basket $x_1 = (0.054608356, 0.004879507)$, which resulted for the MF in the experiment under $p_1 = (1, 12.80710378)$, versus basket $x_2 = (0.040385412, 0.004015199)$, which resulted under $p_2 = (1, 6.958750971)$. To uniquely identify a Walrasian budget line, it is sufficient to have information about its slope and a point lying on the line. The slope corresponds to the observed price ratio (exchange rate) between the two goods traded. The resulting basket (x_1 or x_2, respectively) provides the desired point. The corresponding budget equations describing affordable combinations of the two goods (P, C) under the two prices are thus given by:

$$\begin{aligned}\text{Basket } x_1 : C &\leq (-1/12.80710378) * (P - 0.054608356) \\ &\quad + 0.004879507 \\ \text{Basket } x_2 : C &\leq (-1/6.958750971) * (P - 0.040385412) \\ &\quad + 0.004015199.\end{aligned}$$

Is basket x_2 affordable under price p_1? Total expenditure for x_2 under p_1 would be (Table 1):

$$\begin{aligned}&0.040385412 + (0.004015199 \times 12.80710378) \\ &\quad = 0.091805517 < 0.117093709.\end{aligned}$$

Thus, since x_2 was affordable, but x_1 was chosen under p_1, we have: $x_1 \succeq_D x_2$. On the other hand, expenditure for x_1 under p_2 would be:

$$\begin{aligned}&0.054608356 + (0.004879507 \times 6.958750971) \\ &\quad = 0.088561317 > 0.068325797.\end{aligned}$$

Thus, since x_1 was not affordable when x_2 was chosen, we have: not $x_2 \succeq_D x_1$. We conclude that WARP holds between baskets x_1 and x_2. The revealed preference analysis can also be seen diagrammatically, by using the actual baskets above (Fig. 3). The blue budget set indicates affordable choices for the MF under price vector $p_1 = (1, 12.80710378)$, where the price of P is normalized to 1, and the price of C is 12.80710378, measured in biomass per 1 nmol of P. The resulting consumption vector of P and C, respectively, is $x_1 = (0.054608356, 0.004879507)$. Notice that basket x_2 lies within the blue budget set and, hence, is affordable under p_1. Similarly, the red budget set indicates affordable choices for the MF under price vector $p_2 = (1, 6.958750971)$. The resulting consumption vector of P and C, respectively, is $x_2 = (0.040385412, 0.004015199)$. Notice that basket x_1 lies outside the red budget set and, hence, is unaffordable under p_2.

In *Materials and Methods*, we demonstrate that GARP holds in both markets and under all available data combinations. The resulting consistent rankings are $x_1 \succ x_2 \succ x_3$ for market 1 and $x_4 \succ x_5 \succ x_6$ for market 2. Hence, MF are trading consistently with WE in both markets. Given unlimited data, GARP constitutes a necessary and sufficient condition for WE behavior. However, it may not be as sensitive when, for example, there is little variation in relative prices or quantities. While the amount of available data in the experiment we used is finite, it is worth mentioning that the data contain significant variation. The rich MF treatment involves a rate of change from lowest to highest price of 148%, while the poor MF treatment involves a corresponding 73% price increase. The traded quantities involved also include substantial variation.

Discussion

The Walrasian general equilibrium model is the cornerstone of the economic analysis of markets. This study formally investigates the applicability of the model to studying outcomes in biological markets. Kummel and Salant (2006) (34) studied a theoretical model of a biological market in which a single plant can be supplied with nitrogen by several heterogeneous mycorrhizal trading partners. Their model asserts that, at the optimum, the plant's marginal cost should be equalized across the plant's suppliers. While their model concentrates on the behavior of the plant, the marginal cost-equalization principle might apply more generally. Our starting point consisted of a set of experimental observations of the aggregate transfers of P from the fungal network to the host roots. We focused on whether these transfers, together with the implied exchange rates, satisfy the axioms of revealed preference theory. While undoubtedly there exists some heterogeneity within each MF patch, when we mapped the experimental data to the GET framework, we treated the MF within each patch (poor and rich, respectively) as homogenous. In addition, we considered the total endowment of P within each patch to be the one after any intranetwork transfers. We showed that in each market, the resulting baskets, together with the implied exchange rates, satisfy the axioms of revealed preference, a sufficient condition for WE behavior.

WE is a cornerstone in economics, and its properties have been studied by economists for over 150 y. It is a powerful modeling tool that generates a variety of measurable predictions. While WE describes outcomes in the limiting case of perfect competition, its prescriptions can serve as a benchmark for trades where these conditions are met only approximately. If it proves to be applicable in the study of biological markets, the large body

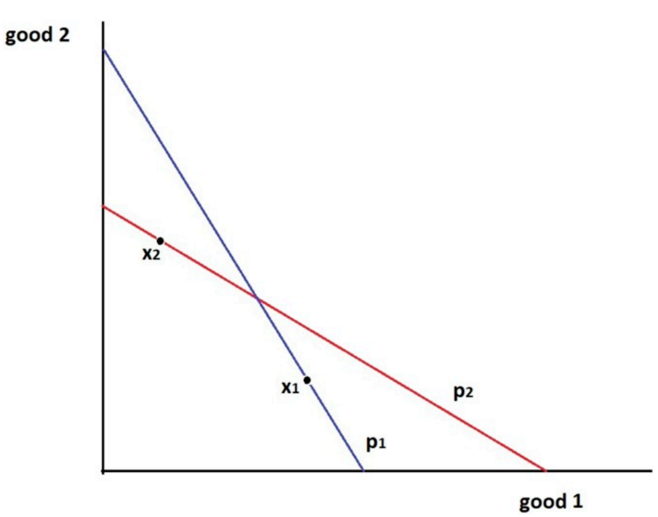

Fig. 2. Failure of WARP.

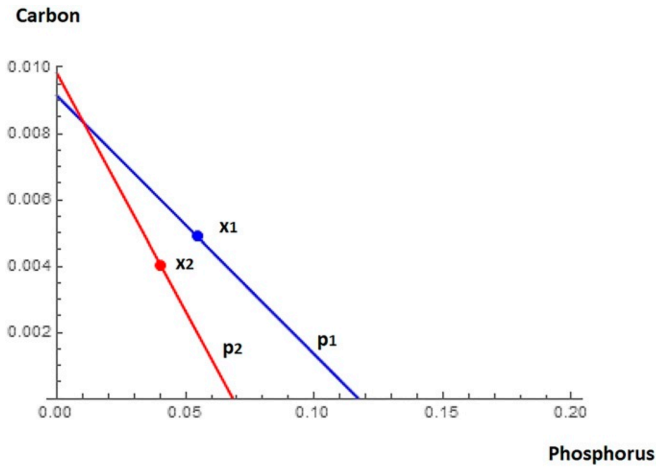

Fig. 3. WARP.

of theoretical results regarding existence, multiplicity, comparative statics, dynamics, and welfare properties of WE can provide additional technical tools for understanding biological phenomena. Part of the power in the axiomatic treatment of WE stems from its ability to make predictions that are independent of the details of how transactions take place. This is in contrast to the game-theoretic models often used in current biological studies. The set of equilibria in game-theoretic models of trade tends to depend on the details of the assumed noncooperative protocol. In many cases, where such details are nonobservable to the modeler, WE can still serve as a useful benchmark. Much like in models of economic behavior involving human participants, we expect that there are examples of biological trade where strategic considerations are sufficiently subdued for the GET model to be useful, and yet others where they are sufficiently dominant for game-theoretic models to be the appropriate tool. Furthermore, there are well-known examples involving human traders where the predictions of WE and those of game-theoretic analysis coincide. At the very least, the techniques outlined in this paper allow researchers to rigorously examine the validity of the Walrasian hypothesis for the specific biological market under study.

The analysis could be extended in many ways. Generating more data would build further confidence in our revealed preference conclusions. This analysis concentrated on the trading behavior of the MF. Future research focusing on the hosts' trading can assist in understanding the plant roots' market behavior. So far, direct experimental measures of C flows have proved challenging. New experiments that measure these flows directly would provide useful data for a more detailed analysis of this biological market [see also Schroeder et al. (5)].

It is also worth mentioning that MF–plant networks can include trade in additional resources, including N and H_2O. It would be interesting to examine the Walrasian hypothesis in the context of trading in multiple goods. MF–plant interactions form only one example of a biological market. The methodology in this paper will hopefully encourage additional experiments that will study the Walrasian hypothesis in other biological markets. Lastly, biological markets provide a promising laboratory for the study of a number of interesting economic questions, including comparative statics, multiple equilibria, equilibrium stability, and the connection between Walrasian and Nash equilibrium.

Materials and Methods

The General Economic Equilibrium Model. A *pure exchange economy* is defined as a tuple $\mathcal{E} = \left\langle I, X^i, \succeq^i, w^i \right\rangle$. I is the set of trading agents; in our case, MF and plant roots. X is the *consumption possibility set*, indicating the combinations of commodities to be evaluated by the traders. In our case, traders evaluate nonnegative amounts of baskets consisting of combinations of two goods: P and C. Thus, $X = \mathbb{R}^2_+$. Traders' preferences are described by a binary preference relation, $\succeq^i$, where $x \succeq^i y$ reads: "Commodity combination x is considered at least as good as combination y by trader i." For all traders i, the relation $\succeq^i$ is usually assumed to be reflexive ($x \succeq^i x$), transitive (for all x, y, z in X, $x \succeq^i y$ and $y \succeq^i z$ implies $x \succeq^i z$) and complete (for all x, y in X, $x \succeq^i y$ or $y \succeq^i z$). Often, preferences are not observed directly. Revealed preference theory allows us to make inferences about preferences from observed trading choices. Finally, w^i stands for the *endowment vector* of trader i. This gives the amount of each good that each trader has access to prior to the commencement of trade. An *allocation* $x = \{x^i\}_{i \in I}$ describes the final consumption quantities of the two goods (P and C) allocated after trade to each of the two types of trading agents (MF and plants). A WE consists of an allocation and prices that satisfy two requirements. First, taking prices and endowments as given, each trader chooses the trades and implied consumption baskets that maximize her payoff subject to her budget (affordability) constraint. The latter requires that the value of her total consumption cannot exceed the value of her initial endowment. Second, the resulting allocation must be feasible. The overall (aggregate) feasibility of the WE allocation is analogous to a conservation law. It requires that the total amount of each good allocated across all market participants after trade does not exceed the total initial endowment of each good. In other words, goods are neither created nor destroyed during the trading process. This requirement is reasonable for both human and biological trading. Formally:

Definition 1. A WE for $\mathcal{E}$ is an allocation $(\widehat{x}^i)_i$ and a price vector $\widehat{p}$ satisfying the following conditions: 1) For all i, $\widehat{x}^i$ is $\succeq^i$-maximal for trader i among those satisfying $p \cdot x^i \leq p \cdot w^i$; and 2) the allocation $(\widehat{x}^i)_{i \in I}$ is feasible—i.e., for every good j, we have: $\sum_{i \in I} x^i_j \leq \sum_{i \in I} w^i_j$.

Definition 2. An allocation $\bar{x}$ is Pareto-efficient if it is feasible and there does not exist another feasible allocation that is preferred to $\bar{x}$ by all traders.

The First Fundamental Theorem of Welfare Economics asserts that if traders' preferences are not saturated, a WE allocation is Pareto-efficient. In other words, while motivated only by selfish considerations, WE behavior implies an efficient level of consumption, where no resources are wasted.

Data. Table 3 is derived by using data from Whiteside et al. (23) on the amount of P traded, from their figure 2A. Data on P retained come from their figure 2B. The prices p_i are derived by exponentiating the natural logarithm of the exchange rates reported in their figure 4C. We use their point estimates, although they also report CIs around these values. These

Table 3. Prices and Consumption Baskets

Basket/ineq	P Sold	P Kept	Price p_i	C Bought	Expend.
Rich					
x_1-None	0.000381	0.054608356	12.80710378	0.004879507	0.117100703
x_2-Med.	0.000577	0.040385412	6.958750971	0.004015199	0.068326184
x_3-High	0.000783	0.010051909	5.155169512	0.004036498	0.030860739
Poor					
x_4-None	0.000343	0.056380193	14.2962891	0.004903627	0.126483865
x_5-Med.	0.000312	0.033355445	17.46152694	0.005447996	0.128485781
x_6-High	0.000237	0.006889843	24.77908622	0.005872643	0.152408581

https://doi.org/10.1073/pnas.2020961118

prices, together with the values of P sold, allow us to infer the amounts of C received (in units of mass obtained by the MF per unit of P transferred). Total expenditure of a given basket is determined by multiplying the quantity consumed by the respective price.

Revealed Preference. Here, we perform a revealed preference analysis. This asserts that the baskets chosen under the three different scenarios (high-/medium-/no-inequality) for each market (rich–poor) are consistent with GARP. Thus, MF trading is consistent with WE behavior.

Basket 1 versus basket 2. Expenditure of x_2 under p_1: $0.091805517 < 0.117093709$. Thus, $x_1 \succ x_2$.

Expenditure of x_1 under p_2: $0.088561317 > 0.068325797$. Thus, not $x_2 \succ x_1$.

Basket 1 versus basket 3. Expenditure of x_3 under p_1: $0.061740961 < 0.117093709$. Thus, $x_1 \succ x_3$.

Expenditure of x_1 under p_3: $0.079759601 > 0.030857489$. Thus, not $x_3 \succ x_1$.

Basket 2 versus basket 3. Expenditure of x_3 under p_2: $0.038138433 < 0.068325797$. Thus, $x_2 \succ x_3$.

Expenditure of x_2 under p_3: $0.061082737 > 0.030857489$. Thus, not $x_3 \succ x_2$. The resulting GARP-consistent ranking for Rich MF is: $x_1 \succ x_2 \succ x_3$.

Basket 4 versus basket 5. Expenditure of x_5 under p_4: $0.111240053 < 0.126487777$. Thus, $x_4 \succ x_5$.

Expenditure of x_4 under p_5: $0.142013841 > 0.128488421$. Thus, not $x_5 \succ x_4$.

Basket 4 versus basket 6. Expenditure of x_6 under p_4: $0.090846851 < 0.126487777$. Thus, $x_4 \succ x_6$.

Expenditure of x_4 under p_6: $0.177896409 > 0.15241691$. Thus, not $x_6 \succ x_4$.

Basket 5 versus basket 6. Expenditure of x_6 under p_5: $0.109444169 < 0.128488421$. Thus, $x_5 \succ x_6$.

Expenditure of x_5 under p_6: $0.168351437 > 0.15241691$. Thus, not $x_6 \succ x_5$ The resulting GARP-consistent ranking for poor MF is: $x_4 \succ x_5 \succ x_6$.

Data Availability. All study data are included in the article and/or supporting information. All experimental data used in this paper are taken from the Whiteside at al. (2019) study and are publicly available at GitHub (https://github.com/gijsbertwerner/Mycorrhizal_inequality) (34).

ACKNOWLEDGMENTS. I thank R. Ford Denison for pointing me to the work by Toby Kiers' group, Ronald Noë for suggesting related references, and Andy Dobson for comments. Pierre Loch provided computational assistance. This research started while I was a visiting scholar at the Institute for the Environment and Sustainability at the University of California Los Angeles and was completed while I was a Council on Foreign Relations Fellow in International Economics. I am grateful to both institutions and their staff for their hospitality. I thank two anonymous referees for several insightful comments. I acknowledge the influence of Leo Hurwicz, 1917–2008, who patiently taught me the logic and power of WE.

1. R. Noë, P. Hammerstein, Biological markets. *Trends Ecol. Evol.* **10**, 336 (1995).
2. G. D. A. Werner *et al.*, Evolution of microbial markets. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 1237–1244 (2014).
3. T. R. Horton, *Mycorrhizal Networks* (Ecological Studies, Springer, Dordrecht, Netherlands, 2015), vol. 224.
4. R. Noë, E. T. Kiers, Mycorrhizal markets, firms, and co-ops, *Trends Ecol. Evol.* **33**, 10 (2018).
5. J. W. Schroeder *et al.*, Mutualist and pathogen traits interact to affect plant community structure in a spatially explicit model. *Nat. Commun.* **11**, 2204 (2020).
6. M. J. Smith, *On Evolution* (Edinburgh University Press, Edinburgh, Scotland, 1972).
7. E. Akcay, J. Roughgarden, Negotiation of mutualism: Rhizobia and legumes. *Proc. Roy. Soc. Lond. B* **274**, 25–32 (2007).
8. E. Sheffer, S. A. Batterman, S. A. Levin, L. O. Hedin, Biome-scale nitrogen fixation strategies selected by climatic constraints on nitrogen cycle. *Nature Plants*, 15182 (2015).
9. S. Bowles, P. Hammerstein, "Does market theory apply to biology?" in *Genetic and Cultural Evolution of Cooperation*, P. Hammerstein, Ed. (MIT Press, Cambridge, MA, 2003), pp. 153–165.
10. P. Hammerstein, R. Noe, Biological trade and markets. *Phil. Trans. R. Soc. B* **371**, 20150101 (2016).
11. R. Noë, How do biological markets compare to the markets of economics? Munich Personal RePEc Archive [Preprint] (2016). https://mpra.ub.uni-muenchen.de/72509/ (Accessed 3 October 2019).
12. A. Bshary, R. Noë, "Biological markets: The ubiquitous influence of partner choice on the dynamics of cleaner fish-client reef fish interactions" in *Genetic and Cultural Evolution of Cooperation*, P. Hammerstein, Ed. (MIT Press, Cambridge, MA, 2003), pp. 167–184.
13. R. F. Denison, E. Toby Kiers, Life histories of symbiotic rhizobia and mycorrhizal fungi. *Curr. Biol.* **21**, R775–R785 (2011).
14. E. C. Hammer, J. Pallon, H. Wallander, P. Olsson, Tit for tat? A mycorrhizal fungus accumulates phosphorus under low plant carbon availability. *FEMS Microbiology Ecology* **76**, 236–244 (2011).
15. E. T. Kiers, R. A., Rousseau, S. A. West, R. F. Denison, Host sanctions and the legume-rhizobium mutualism. *Nature* **425**, 78–81 (2003).
16. E. T. Kiers, R. A. Rousseau, R. F. Denison, Measured sanctions: Legume hosts detect quantitative variation in rhizobium cooperation and punish accordingly. *Evol. Ecol. Res.* **8**, 1077–1086 (2006).
17. E. T. Kiers *et al.*, Reciprocal rewards stabilize cooperation in the mycorrhizal symbiosis. *Science* **12**, 880–882 (2011).
18. E. G. Weyl, M. E. Frederickson, D. W. Yu, N. E. Pierce, Economic contract theory tests models of mutualism. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 15712–15716 (2010).
19. G. Debreu, *Theory of Value: An Axiomatic Analysis of Economic Equilibrium* (Yale University Press, New Haven, CT, 1959).
20. K. J. Arrow, F. H. Hahn, *General Competitive Analysis* (Advanced Textbooks in Economics, Elsevier Science, Amsterdam, Netherlands, 1971), vol. 12.
21. D. Gale, Bargaining and competition. Part I: Characterization. *Econometrica* **54**, 785–806 (1986).
22. B. Andreas, P. Martin, The most widespread symbiosis on Earth. *PLoS Biol.* **4**, e239 (2006).
23. M. D. Whiteside *et al.*, Mycorrhizal fungi respond to resource inequality by moving phosphorus from rich to poor patches across networks. *Curr. Biol.* **29**, 2043–2050.e8 (2019).
24. A. A. Alchian, Uncertainty evolution, and economic theory. *J. Polit. Econ.* **57**, 211–221 (1950).
25. M. Schaffer, Evolutionarily stable strategies for a finite population and a variable contest size. *J. Theor. Biol.* **132**, 469–478 (1988).
26. F. Vega-Redondo, The evolution of Walrasian behavior. *Econometrica* **65**, 375–384 (1997).
27. C. Kim, K. C. Wong, Evolution of Walrasian equilibrium in an exchange economy. *J. Evol. Econ.* **21**, 619–647 (2011).
28. S. Gjerstad, J. Dickhaut, Price formation in double auctions. *Game. Econ. Behav.* **22**, 1–29 (1998).
29. L. Hurwicz, R. Radner, S. Reiter, A stochastic decentralized resource allocation process: Part I. *Econometrica* **43**, 363–393 (1975).
30. P. A. Samuelson, A note on the pure theory of consumers' behaviour. *Economica* **5**, 61–71 (1938).
31. D. J. Brown, R. L. Matzkin, "Testable restrictions on the equilibrium manifold" in *Computational Aspects of General Equilibrium Theory. Refutable Theories of Value.* D. Brown, F. Kubler, Eds. (Lecture Notes in Economics and Mathematical Systems, Springer, Berlin, Germany, 2008), vol. 604, pp. 11–25.
32. A. Fostel, H. E. Scarf, M. J. Todd, Two new proofs of Afriat's theorem. *Econ. Theor.* **24**, 211–219 (2004).
33. M. Kummel, S. W. Salant, The economics of mutualisms: Optimal utilization of mycorrhizal mutualistic partners by plants. *Ecology* **87**, 892–902 (2006).
34. G. Werner, Mycorrhizal_inequality. GitHub. https://github.com/gijsbertwerner/Mycorrhizal_inequality. Deposited 27 April 2019.