Generalizing the Proportional Response Dynamic for Exchange Economies

Cooper Simpson[1]

[1]University of Colorado Boulder
cooper.simpson@colorado.edu

December 4, 2023

Abstract

We investigate generalizing the proportional response dynamic for exchange economies via two extensions: arbitrary endowments and arbitrary network structures. In a complete graph with arbitrary endowments, we show that a simple strategy of splitting up bids equally between sellers of identical goods is a Nash equilibrium. Further, we argue that this strategy is the natural approach for choosing initial bids, implying this is the ideal approach in general. For an arbitrary network with arbitrary endowments, we propose a method for updating bids based on the individual utility extracted for each version of a good. We also prove that these two strategies are identical in a complete graph under a mild assumption on the initial bids.

1 Introduction

In this work we will consider modeling the dynamics of exchange economies – an economic setting where a group of agents trades a selection of commodities amongst themselves. It has been shown in a plethora of work that market equilibria exist under mild conditions in a number of variations within this general setting. In particular, it can be shown that equilibria exist even when the network structure describing agent interactions is arbitrary. However, establishing reasonable dynamics with strong convergence guarantees has been difficult. The Proportional Response dynamic is a simple mechanic which has been recently investigated in this regime, and shown to have promising convergence characteristics. Despite this, there are a few key limitations to the existing framework, which we will address here. In particular, we will start by trying to generalize PR to arbitrary agent endowments. We will then consider moving beyond a complete graph to an arbitrary network. Both the details of our generalizations to PR and their convergence properties will be discussed.

1.1 Related Work

The ideas surrounding economic market equillibria have been around since [Walras, 1896], and this includes associated dynamics. [Arrow and Debreu, 1954] provided the first solid existence results for exchange economies (in the complete case), which has been thoroughly expanded upon since. Notable to the results here, is the work of [Kakade et al., 2004], where they generalize the exchange economy to an arbitrary network setting, and provide equilibria existence results similar to Arrow and Debreu. More recently, the work of [Andrade et al., 2021] has fixed a few issues in this generalized setting while maintaining strong equilibrium results.

A large amount research has been put towards computing equilibria, and in some sense this can be seen as a dynamic. For example, casting it is an optimization problem, one can view the trajectory of an optimizer as the dynamics of the market. However, the unfolding of events under that lens does not necessarily have any connection to the "real world". Ideally, the dynamics would be some sort of natural interaction between agents, that more effectively captures the behaviour of real markets. [Zhang, 2011] introduces the Proportional Response (PR) dynamic, although only in the context of Fisher markets – an exchange economy with distinct buyers and sellers. Really, this is an extension of the trading post mechanism of [Shapley and Shubik, 1977], which allocates goods in proportion to bids. It is the mechanism for updating bids in proportion to utility received from those goods that makes PR its own dynamic.

The main work we build off of here is that of [Brânzei et al., 2021] and the follow up in [Brânzei, 2021]. The PR dynamic is extended to the more general exchange economy setting with linear utilities. This dynamic is shown to converge to market equilibrium, with the possibility of cycling prices (albeit in a well defined manner). Important to this work is the fact that only one unique good per agent is considered. At equilibrium this can be shown to be equivalent to arbitrary endowments, but the dynamics along the way may still be quite different.

1.2 Our Contributions

We will present two extensions to the PR dynamic that attempt to generalize it to a broader set of exchange economy settings. First, we will discuss a change to the bid update that allows for arbitrary endowments when the graph is complete. This tweak is justified via a Nash equilibrium result for the optimal strategy in splitting a bid between different versions of a good. We will also argue that this strategy would arise naturally. The second change we will propose is a bid update rule that allows for arbitrary endowments in an arbitrary graph. We will also show that in a complete graph, the weighted split strategy converges to this latter strategy. We then discuss the convergence characteristics of this new dynamic.

2 Preliminaries

Before discussing the generalizations to the Proportional Response (PR) dynamic that this work proposes, we must first establish some important notation. This is similar to that of [Brânzei et al., 2021] and [Brânzei, 2021], but is somewhat more general to facilitate the extensions discussed in later sections. We will also review the PR dynamic itself, and its associated convergence guarantees. In the exchange economy setting we consider n agents and m divisible commodities. We view the underlying structure of this economy as an undirected graph, so we denote the neighbor set of agent i as $N(i)$. An edge between two agents indicates the ability for them to trade. It is perhaps important to point out that $N(i)$ could contain i, indicating that an agent could "trade" with themselves.

Each agent $i \in [n]$ has an endowment of goods $\mathbf{e}^i \in \mathbb{R}^m_+$, and a linear utility for those goods $\mathbf{u}^i \in \mathbb{R}^m_+$. On this latter point, we note that such an assumption may be somewhat unreasonable, but it is still an important case which is not necessarily easier to handle. We will consider a discrete dynamical setting, so we denote these time steps as $t \in \mathbb{Z}_+$. Then for $j \in N(i)$ we define the following: agent i's bid on agent j's goods as $\mathbf{b}^{ij}(t) \in \mathbb{R}^m_+$, and the allocation of agent j's goods received by agent i as $\mathbf{x}^{ij}(t) \in \mathbb{R}^m_+$. The prices of an agent's goods can be viewed as the sum of all bids on those goods, $\mathbf{p}^i(t) = \sum_j \mathbf{b}^{ji}$. Lastly, we denote agent i's budget available in the next round as $B_i(t+1) = \sum_k \mathbf{p}^i_k(t)$, where $B_i(0)$ is initialized in some appropriate manner. All of this notation is summarized in the following table:

Quantity	Notation	Details
Neighbor Set	$N(i)$	Predetermined
Endowment	$\mathbf{e}^i$	Predetermined
Utility	$\mathbf{u}^i$	Predetermined
Bid	$\mathbf{b}^{ij}(t)$	Computed
Allocation	$\mathbf{x}^{ij}(t)$	Computed
Price	$\mathbf{p}^i(t)$	Defined as $\sum_j \mathbf{b}^{ji}$
Budget	$B_i(t+1)$	Defined as $\sum_k \mathbf{p}^i_k(t)$

Table 1: Proportional Response dynamic notation

2.1 Proportional Response Dynamic

The PR dynamic operates in the setting of a complete graph. We will assume that each agent is endowed with a single unit of a unique good, i.e. $\mathbf{e}^i_k = \mathcal{I}_{k=i}$, and if $\mathbf{u}^i_k > 0$ then $\mathbf{b}^{ik}_k(0) > 0$. A generalized form of the dynamic from [Brânzei et al., 2021] is given as follows:

Definition 1: Proportional Response

For each $t \in \mathbb{Z}_+$, compute the following:

1. **Exchange goods:**
$$\mathbf{x}^{ij}_k(t) = \frac{\mathbf{b}^{ij}_k}{\sum_i \mathbf{b}^{ij}_k} \mathbf{e}^i_k$$
2. **Compute utility:**
$$U_i(t) = \sum_j \langle \mathbf{u}^i, \mathbf{x}^{ij}(t) \rangle$$
3. **Update bids:**
$$\mathbf{b}^{ij}_k(t+1) = \frac{\langle \mathbf{u}^i, \mathbf{x}^{ij}(t) \rangle}{U_i(t)} B_i(t+1)$$

In fact, many of the computations above simplify further due to the sparsity of endowments. For example, the inner product in the second step measures the total utility extracted from another agent, which simplifies to $\mathbf{u}^i_j \mathbf{x}^{ij}_j$. We state it in this manner to better connect it with the generalizations we present latter.

We have not yet made clear what constitutes a market equilibrium; definition 2 remedies this, and it provides us with a general statement that applies to arbitrary networks.

Definition 2: Graphical Equilibrium

In a graphical exchange economy, an equilibrium is a set of prices $\hat{\mathbf{p}}^i$ and allocations $\hat{\mathbf{x}}^{ij}$ that satisfy the following:

- **Market Clearing:** $\sum_{i \in N(j)} \hat{\mathbf{x}}^{ij} = \mathbf{e}^j, \ \forall j \in [n]$
- **Optimality:** For all $i \in [n]$, the allocations $\hat{\mathbf{x}}^{ij}$ maximize U_i with respect to all allocations that satisfy $\sum_{j \in N(i)} \langle \hat{\mathbf{p}}^j, \mathbf{x}^{ij} \rangle \leq \langle \hat{\mathbf{p}}^i, \mathbf{e}^i \rangle$.

In [Brânzei et al., 2021], the authors consider a slightly more general version of PR than we have discussed so far. Namely, they allow for a "lazy" version of the dynamic, where agents save a fraction $\alpha \in [0, 1)$ of their wealth at each round. They then show that any fixed point of this dynamic is a market equilibrium. Form here, it is shown that each agent's utility converges to the market equilibrium, and each agent's allocation converges to a market equilibrium. Note the difference in qualifiers, as equilibrium utilities are necessarily unique, but allocations are not. When it comes to prices, it is shown that they converge to a market equilibrium for $\alpha \neq 0$. If this is not the case, then it is possible for them to cycle, but in a well characterized manner.

3 Generalizing Endowments

The most immediate question arising from what we have discussed so far about the Proportional Response (PR) dynamic, and one that is posed in [Brânzei et al., 2021] and [Brânzei, 2021], is that of how to move beyond one good per agent? The issue lies in determining how an agent should choose to split up a bid between all available versions of a good. At its core, this is a question of how the agent values these versions, including its own. Many different approaches could be taken to modeling this situation appropriately, but we will make the assumption that an agent's utility for a good is independent of the seller of that version. We do not claim this is the ideal approach, but it allows us to focus on the dynamics.

Let's define the total bid available to agent $i \in [n]$ for each good $k \in [m]$ as $\mathbf{b}^i(t) \in \mathbb{R}^m_+$. Initialize this in the same way you would for the standard PR dynamic, and for $t \in \mathbb{Z}_+$ compute it as follows:

$$\mathbf{b}^i_k(t+1) = \frac{\mathbf{u}^i_k \sum_l \mathbf{x}^{il}_k}{U_i(t)} B_i(t+1) \tag{1}$$

which is the PR bid update for a unique good. The following theorem outlines our proposed strategy for splitting up a bid for a given good amongst available sellers of that good:

Theorem 1

In a complete graph with arbitrary endowments, the following strategy for splitting up a bid is a Nash equilibrium:

$$\mathbf{b}^{ij}_k = \frac{\mathbf{e}^j_k}{\sum_l \mathbf{e}^l_k} \mathbf{b}^i_k \tag{2}$$

in the sense that unilateral deviation decreases an agent's utility.

To prove this result, we will assume one agent has deviated from the proposed strategy, and show that their utility can be increased by moving back towards the weighted split.

Proof

Fix a particular agent and commodity, where we denote this agent's budget for this good as B. Let b_i denote the bid by this agent on agent i's version of this good, let e_i denote agent i's endowment of this good, and define E_i as follows:

$$E_i = \frac{e_i}{\sum_l e_l}$$

The case where there is only one seller of the good in question is trivial, so we will assume that there are at least two. Furthermore, we will assume that there are at least two agents with positive utility for this good, and thus have positive bids associated with it. If this is not the case, then any split which ensures a positive bid on all versions of a good is globally optimal, so the result still holds.

Suppose that there exists an $i, j \in [n]$ such that $b_i < E_i B$ and $E_j B < b_j$, so the bid is not split according to the Nash strategy. Define $\tilde{b}_i = b_i + \epsilon$, $\tilde{b}_j = b_j - \epsilon$, and $\tilde{b}_l = b_l$ for $l \neq i, j$; where $\epsilon > 0$ is sufficiently small enough to maintain the budget constraint. Define the allocation of the good in question as follows:

$$f(\epsilon) = \sum_l \frac{\tilde{b}_l}{\tilde{b}_l + a_l} e_l$$

The quantity $a_l = cE_l$, for $c > 0$ some constant, is the total amount bid by other agents on agent l's version of the good. It has this form, because we assume that all other agents are operating under the Nash strategy. Now, observe the following:

$$\frac{d}{d\epsilon} f(\epsilon) = \frac{a_i e_i}{(\tilde{b}_i + a_i)^2} + \frac{a_j e_j}{(\tilde{b}_j + a_j)^2} = \frac{(b_j - \epsilon + a_j)^2 a_i e_i - (b_i + \epsilon + a_I)^2 a_j e_j}{(b_i + \epsilon + a_i)^2 (b_j - \epsilon + a_j)^2}$$

Choosing $\epsilon < \min\{b_j - E_j B, E_i B - b_i\}$ ensures the following inequalities hold:

$$b_j e_i + a_j e_i > \frac{e_j e_i}{\sum_l e_l} B + \frac{e_j e_i}{\sum_l e_l} c + \epsilon e_i$$
$$b_i e_j + a_i e_j < \frac{e_i e_j}{\sum_l e_l} B + \frac{e_i e_j}{\sum_l e_l} c - \epsilon e_j$$

Combining these with a few algebraic manipulations results in the following:

$$(b_j - \epsilon + a_j)^2 a_i e_i > (b_i + \epsilon + a_i)^2 a_j e_j$$

Thus, we conclude that $f'(\epsilon) > 0$ on $(0, \epsilon)$, which implies the allocation is increasing on that interval. Therefore, any deviation from the Nash strategy reduces the allocation an agent receives, and because agents only bid on goods which provide positive utility, a reduction in allocation corresponds to a reduction in utility. ▲

We can also argue that this strategy is not only some strange isolated equilibrium, but is, in fact, a reasonable strategy to arise naturally. Consider an agent showing up to the market unaware of what other agents are bidding on. Without this knowledge, the best action the agent can take is to at least bid some positive amount on each seller of a good for which they have utility. This ensures that the agent receives a maximal allocation in the case where no other agents desire the same goods. From here, it can only be better to use the weighted split strategy, because if there turn out to be other agents bidding on the same goods, this is then optimal.

4 Generalizing Network Structure

As was shown in [Kakade et al., 2004], the generalization to arbitrary network structures in the setting of exchange economies is quite powerful. It allows one to incorporate more local interactions to yield phenomenon not present in the complete case, and, importantly, strong equilibria guarantees still hold. Thus, the same motivation for investigating dynamics applies here, and in particular, we want to extend the Proportional Response (PR) dynamic with arbitrary endowments to this setting.

A reasonable question to ask is why weighted splitting is no longer a valid approach? Consider the following incomplete graph, where G is some arbitrary sub-graph:

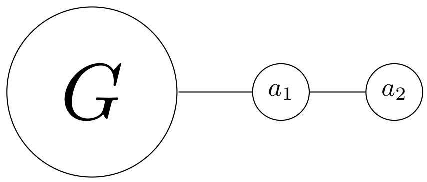

Suppose that a_2 has a good for which a_1 has utility, and further suppose that the same good can be obtained by a_1 in G. In this scenario, a_1 can obtain a_2's entire endowment for effectively a price of 0, because there is no competition. A bid of any $\epsilon > 0$ by a_1 on the version of the good in question would guarantee the full endowment. Thus, it is beneficial for a_1 to move more of their bid for the good in question to sellers in G, so we see that weighted splitting is no longer a Nash equilibrium.

Standard PR updates bids in proportion to the utility extracted from a good relative to the total utility extracted by that agent. Given this, it seems the natural extension is to perform this same process for each individual version of a good, so we propose the following bid update rule:

$$\mathbf{b}_k^{ij}(t+1) = \frac{\mathbf{u}_k^i \mathbf{x}_k^{ij}}{U_i(t)} B_i(t+1) \tag{3}$$

where for $t = 0$, and for any given agent, a positive bid is made on each version of a good available to said agent when that has utility for that good, and a bid of zero is made otherwise.

This doesn't necessarily address the issue raised with our example at the beginning of this section. However, the overall question becomes much more complicated as we must factor in new details, such as whether an agent is aware of the local (or any level) structure of the network. Equation (3) makes much more sense if they are not aware of their neighbors connections than if they are.

4.1 Connection To Weighted Split

A reasonable question to ask regarding the bid update rule of Eq. (3) is how it compares to the weighted split in the complete case. The following lemma makes this connection clear:

Lemma 1

In a complete graph, the weighted split bid update rule of Eq. (2) is equivalent to the per-version bid update rule of Eq. (3), assuming the previous bid was computed according to the weighted split.

Proof

Recall the weighted split bid update:

$$\mathbf{b}_k^{ij} = \frac{\mathbf{e}_k^j}{\sum_l \mathbf{e}_k^l}\mathbf{b}_k^i$$

where we have removed the dependence on t as that is arbitrary. Using Eq. (1), we can rewrite this as follows:

$$\mathbf{b}_k^{ij} = \frac{\mathbf{e}_k^j}{\sum_l \mathbf{e}_k^l}\frac{\mathbf{u}_k^i \sum_l \mathbf{x}_k^{il}}{U_i}B_i$$

Thus, let us consider the following:

$$\frac{\mathbf{e}_k^j}{\sum_l \mathbf{e}_k^l}\sum_l \mathbf{x}_k^{il} = \frac{\mathbf{e}_k^j}{\sum_l \mathbf{e}_k^l}\sum_l \frac{\frac{\mathbf{e}_k^l}{\sum_r \mathbf{e}_k^r}\mathbf{b}_k^i}{\frac{\mathbf{e}_k^l}{\sum_r \mathbf{e}_k^r}\mathbf{b}_k^i + c_k\frac{\mathbf{e}_k^l}{\sum_r \mathbf{e}_k^r}}\mathbf{e}_k^l$$

which uses the fact that the previous bid was computed according to the weighted split to replace $\mathbf{x}_k^{il}$. The quantity c_k is the is the sum of all other agents' budget available to bid on good k. Some simple rearranging and cancelling of terms yields the equality below:

$$\cdots = \frac{\mathbf{e}_k^j}{\sum_l \mathbf{e}_k^l}\frac{\mathbf{b}_k^i}{\mathbf{b}_k^i + c_k}\sum_l \mathbf{e}_k^l = \frac{\mathbf{b}_k^i}{\mathbf{b}_k^i + c}\mathbf{e}_k^j$$

Now we can multiply and divide the right hand side by $\frac{\mathbf{e}_k^j}{\sum_l \mathbf{e}_k^l}$, which yields the following:

$$\frac{\mathbf{b}_k^i}{\mathbf{b}_k^i + c_k}\mathbf{e}_k^j = \mathbf{x}_k^{ij}$$

Thus, we conclude with the following:

$$\mathbf{b}_k^{ij} = \frac{\mathbf{e}_k^j}{\sum_l \mathbf{e}_k^l}\mathbf{b}_k^i = \frac{\mathbf{u}_k^i \mathbf{x}_k^{ij}}{U_i}B_i$$

so the result holds. ▲

The assumption in Lemma 1 that the previous bid was computed according to the weighted split seems, at first, a bit strange. However, Theorem 1 and the discussion thereafter shows that this a reasonable assumption starting at $t = 2$. The first bid on any given good is positive when an agent has positive utility for that good, and zero otherwise. A rational agent will then use the weighted split to compute the bid updates, and from there, the two updates are equivalent.

In some sense, Lemma 1 gives the convergence of the weighted bid update (Eq. (2)) to the general update (Eq. (3)). It is unclear if the opposite result holds. That is to say, if the general update is employed from $t = 1$, will this strategy eventually converge to the weighted split?

5 Convergence

Our goal is to extend PR to arbitrary endowments, and the the update rule of Eq. (3) accomplishes this in some sense, but the core question is whether or not the new dynamic converges to a market equilibrium. As with standard PR, we can split this into two questions: are fixed points market equilibrium, and does the dynamic necessarily converge? Given the structure of PR, as discussed earlier, the market clearance condition for equilibrium is trivially satisfied and the budget constraint is always saturated. It is also trivial to show that fixed point allocations are optimal with respect to fixed point prices. This should not be all that surprising, as the structure of the PR dynamic enforces the conditions of definition 2 quite rigidly. For completeness, we will rigorously establish this in Lemma 2. Let $\hat{a}$ denote the fixed point value for a quantity a. The fixed point bids can be written as follows:

$$\hat{\mathbf{b}}_k^{ij} = \frac{\mathbf{u}_k^i \hat{\mathbf{x}}_k^{ij}}{\hat{U}_i} \hat{B}_i \tag{4}$$

Lemma 2

A fixed point of the dynamic defined by Eq. (3) is a graphical equilibrium.

Proof

First, we will show that market clearance holds at fixed points. Consider the following:

$$\sum_{i \in N(j)} \hat{\mathbf{x}}_k^{ij} = \sum_{i \in N(j)} \frac{\hat{\mathbf{b}}_k^{ij} \hat{U}_i}{\hat{B}_i \hat{\mathbf{u}_k^i}}$$

which only requires rearranging Eq. (4) and summing. Note that by construction, the following relationship holds for $i \in N(j)$:

$$\hat{\mathbf{b}}_k^{ij} = \frac{\hat{\mathbf{u}_k^i} c_i \mathbf{e}_k^j}{\hat{U}_i} \hat{B}_i$$

where c_i is the fraction of agent j's good k allocated to agent i. This follows directly from the definition of the bid update. Plugging this in gives us the following:

$$\sum_{i \in N(j)} \hat{\mathbf{x}}_k^{ij} = \sum_{i \in N(j)} c_i \mathbf{e}_k^j = \mathbf{e}_k^j$$

which follows because the c_i necessarily sum to one. To see that the fixed point allocations are optimal, one only need consider the possible allocations at the a fixed point. The only way to change the allocation is by changing the associated bids, but because the bids are fixed, there is only one possible allocation. Thus, this single allocation trivially maximizes U_i for all $i \in [n]$. With both conditions of definition 2 satisfied, the result holds. ▲

It remains to be shown that our generalized PR dynamic converges.

Expanding the allocation (definition 1) and budget (Table 1) via definition, we can write the following:

$$\hat{\mathbf{b}}_k^{ij} = \frac{\mathbf{u}_k^i}{\hat{U}_i} \frac{\hat{\mathbf{b}}_k^{ij}}{\sum_l \hat{\mathbf{b}}_k^{lj}} \mathbf{e}_k^j \langle \mathbf{1}, \hat{\mathbf{p}}^i \rangle$$

Assuming the LHS is not zero, we may cancel terms, and noting that the sum in the denominator on the RHS is the price for agent j's good k yields the final result:

$$\hat{U}_i = \frac{\mathbf{u}_k^i \mathbf{e}_k^j \langle \mathbf{1}, \hat{\mathbf{p}}^i \rangle}{\hat{\mathbf{p}}_k^j} \tag{5}$$

6 Discussion

In this work, we have taken a large step towards fully generalizing the Proportional Response dynamic to arbitrary exchange economy settings. In particular we have given modified bid update rules to allow for arbitrary endowments and network structures. However, there is still much work to be done to justify these proposals, and to continue incorporating other ideas.

We have discussed the possibility that fixed points of the weighted splitting dynamic are market equilibria, but we have not guaranteed convergence of the dynamic to these equilibria. In our discussion of the general update rule, we connected it to weighted splitting in the complete case. However, a question this raises, is whether any initialization would converge to this Nash equilibrium? In the incomplete case, as well, it is unclear if this individual bid update converges to market equilibria.

Although PR has appears quite powerful, it is unclear if it is the "correct" dynamic to be considering. Is it possible to show the proportional update this dynamic uses is a Nash equilibrium for splitting up a budget amongst goods? Beyond specific dynamics, it would be interesting to incorporate existing exchange economic models such as resale or production.

References

[1] Gabriel Andrade et al. "Graphical Economies with Resale". In: *Proceedings of the 22nd ACM Conference on Economics and Computation* (July 2021).

[2] Kenneth J. Arrow and Gerard Debreu. "Existence of an Equilibrium for a Competitive Economy". In: *Econometrica* 22.3 (1954), pp. 265–290.

[3] Simina Brânzei. "Exchange markets: proportional response dynamics and beyond". In: *ACM SIGecom Exchanges* 19.2 (2021), pp. 37–45.

[4] Simina Brânzei, Nikhil Devanur, and Yuval Rabani. "Proportional dynamics in exchange economies". In: *Proceedings of the 22nd ACM Conference on Economics and Computation*. 2021, pp. 180–201.

[5] Sham M Kakade, Michael Kearns, and Luis E Ortiz. "Graphical economics". In: *International Conference on Computational Learning Theory*. Springer. 2004, pp. 17–32.

[6] Lloyd Shapley and Martin Shubik. "Trade using one commodity as a means of payment". In: *Journal of political economy* 85.5 (1977), pp. 937–968.

[7] Léon Walras. *Etudes d'économie sociale (Théorie de la répartition de la richesse sociale)*. F. Rouge et cie, sa, 1896.

[8] Li Zhang. "Proportional response dynamics in the Fisher market". In: *Theoretical Computer Science* 412.24 (2011), pp. 2691–2698.