A General Equilibrium Supply and Demand Framework for Demonstrating the Fundamental Theorems of Welfare Economics

March 2021

Luc Noiset
Department of Economics Finance and Quantitative Analysis
Coles College of Business
Kennesaw State University
Kennesaw, Georgia

Abstract

For many decades, the Edgeworth-Bowley box has been the workhorse of the economics profession for illustrating the fundamental theorems of welfare economics. This article presents a device that complements the Edgeworth-Bowley box and its representations of the fundamental theorems. The device is a simplified general-equilibrium model. It is an elaboration of the relatively straightforward concept of maximizing the total social surplus across markets, yet the general equilibrium structure provides completeness that traditional partial equilibrium consumer and producer surplus narratives do not. The model is as accessible as the Edgeworth-Bowley box but provides numerical examples of the fundamental theorems. It can be a useful addition in the toolbox of professors and students of welfare economics but does not appear in popular economic textbooks.

JEL Codes: A22, B41, D61

Key Words: Edgeworth-Bowley Box, Welfare Economics, General Equilibrium

* I thank Nick Baigent, Mark Rider and Jesse Schwartz and reviewers of the Coles Working Paper Series for helpful comments. Remaining errors are mine.

1. Introduction

The first fundamental theorem of welfare economics states that the equilibrium of a perfectly competitive economy will achieve an allocation that is Pareto optimal. The second states that any Pareto optimal allocation can be supported by some competitive equilibrium. These two theorems are arguably the greatest intellectual achievements of the economics profession in the twentieth century, but they are often difficult to illustrate even in advanced undergraduate classrooms.[1, 2]

The famous Edgeworth-Bowley Box is the workhorse of the economics profession that has been traditionally called upon to show that perfect competition achieves Pareto optimal allocations of society's resources. The Edgeworth-Bowley Box highlights the critical tangency conditions of a Pareto optimum in a two-good, two-person, two-input world, and once the tangency conditions are described, the case can be made that these conditions are met in an n-person competitive economy in which all agents are price-takers.[3]

Figure 1 depicts an example of the familiar Edgeworth-Bowley-type exposition for an endowment/exchange economy.[4] The Pareto optimal equilibrium occurs where the indifference curves of the two individuals are tangent to one another and the slope of the tangent, which is the common Marginal Rate of Substitution (MRS), is parallel to the slope of the Marginal Rate of Transformation (MRT) along the Production Possibilities Frontier (PPF).

---

[1] The fundamental theorems of welfare economics were arrived at by the collective intellectual efforts of many economists including Dupuit, Walras, Marshall, Hotelling, Hicks and Pigou to name a few. The modern derivations are generally credited to Arrow and Debreu (1954) and McKenzie (1954).

[2] It should be noted that there are two primary reasons for studying welfare economics. One is to understand the power of Adams Smith's invisible hand conjecture: A competitive equilibrium indeed achieves a Pareto optimum. The second is to be able to reject the "exaggerated claims" that the market economy is "Nirvana on earth." See Stiglitz (1991).

[3] See Humphrey (1996) for a wonderful discussion of the history of the Edgeworth-Bowley box. See also Edgeworth (1881), and Bowley (1924) for a nice historical meander, and Hicks' (1939) defense of welfare economics.

[4] The box diagram using isocosts and isoquants similarly shows the efficient use of two productive inputs across two output markets.

Figure 1

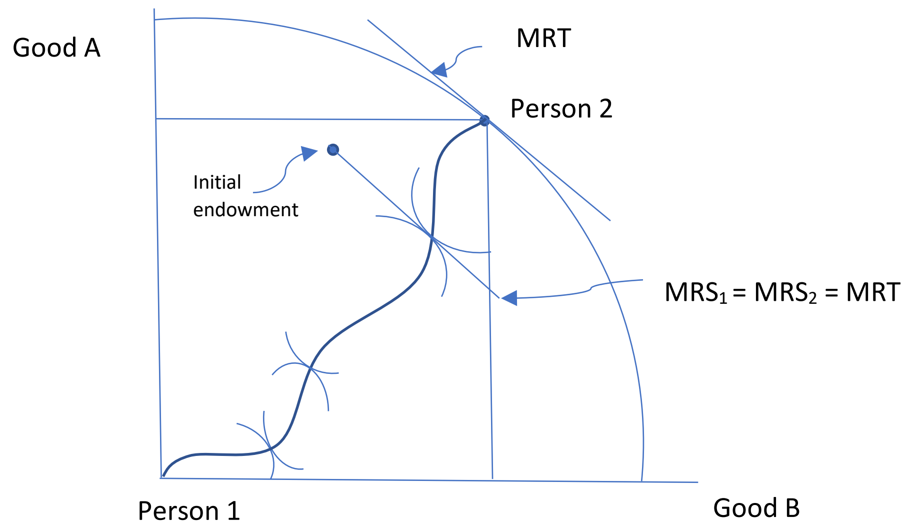

This article presents an alternative pedagogical device that complements the Edgeworth-Bowley-Box representation of the fundamental theorems, but which offers an alternative perspective and a different way to illustrate the standard results. The device is a simple general equilibrium model structured to produce numerical Pareto optimal competitive equilibria arising from a simple demand and supply framework and coinciding with a point on a two-dimensional Production Possibilities Frontier. Most general equilibrium models are too cumbersome for classroom use, but the model presented here is straightforward enough to be suitable for undergraduate microeconomic theory courses. It has the flexibility to address several concepts, especially demonstrating the first and second theorems of welfare economics.

This methodology does not appear to be used in current economics textbooks but can be an enlightening addition to the typical classroom exposition of the fundamental theorems. It also provides a foundation for useful numerical classroom or homework exercises.[5] Basic spreadsheet software can be used to carryout simulations with the model demonstrating computationally, with explicit utility functions, if and in what sense social welfare may or may not be maximized at the perfectly competitive market equilibrium.

The perfect-competition story generates the two key features of Pareto optimality. The first feature, production efficiency, can be conveyed in the classroom in a relatively straightforward way. As the story goes, competition among firms pushes inefficient firms out of business and forces surviving firms to use the best-known production technology to minimize production costs. Meanwhile, flexible prices enable all markets to clear, ensuring that there are no unemployed resources. Together, cost minimization, through the use of the most efficient technology, and market clearing in resource markets, ensure that the economy is operating on its Production Possibilities Frontier.

The second feature of Pareto optimality, allocative efficiency, means that that the economy is producing the correct quantities of each good or service vis-à-vis the quantities of other goods and services produced. Perfect competition generates this result because prices signal the marginal valuations of goods and services while individuals trade to the point where their marginal valuations are equalized.[6] Allocation efficiency is typically the more difficult concept to get across to students. Often, this is attempted within a partial equilibrium framework using consumer and producer surplus concepts with the supply and demand diagram of a single market. A pedagogical technique that is immediately suspect when one considers that it is the quantity of one good relative to the available quantity of other goods that is the relevant concept.

---

[5] Notwithstanding its absence from modern textbooks, this framework has an impressive historical pedigree for examining issues of welfare economics – see for example, Hotelling (1938).
[6] This result is illustrated by the contract curve and the tangencies of indifference curves in equilibrium, the equalization of marginal rates of substitution, in the traditional Edgeworth-Bowley Box exposition.

The model and pedagogical approach proposed in this paper are described more completely in the next section.

2. The Model

The mechanism that allows for this model's simple structure and computational usefulness is the use of the reflection of the demand curve in one market as the marginal social cost curve in the other market. In a perfectly competitive economy with no externalities, the market supply curve is the social marginal cost of producing an incremental unit of a good. This social marginal cost is derived from the opportunity cost of resources needed to produce the incremental unit, and the opportunity cost of resources is equal to the forgone value of the alternative good that could have been produced with those resources. That forgone value is derived from the marginal demand valuations, moving backwards up the demand curve of the alternative good, as incremental units are forgone. In other words, the social marginal cost of resources in one market is measured by the height of the demand curve in the alternative market.[7] Thus, one can derive an *implicit* supply curve in one market by looking at the incremental (*forgone*) demand valuations in the alternative market. That is, the marginal social cost of resources in one market is measured by the height of the demand curve in the alternative market.

The labour-leisure choice problem is suppressed in the model by assuming that each individual owns an endowment, ℓ, of the sole productive resource in the economy, but this resource provides no direct positive or negative utility value. In other words, unlike traditional labour, which can be used for leisure activities, the ℓ in this model has an opportunity cost of zero outside of production. This ensures that all

[7] In other words, an implicit market supply curve, reflecting full social marginal costs, can be derived without reference to the costs or production functions of firms.

productive resources are put to work since it makes no sense for the individual to withhold her $\boldsymbol{\ell}$ from production for any return above zero.[8]

Although there are no production or cost functions needed in the model, production is implied but subsumed by assuming a constant-returns-to-scale (CRS) production structure in which one unit of $\boldsymbol{\ell}$ can be costlessly transformed into one unit of output. Thus, there are no scale economies, the number of firms is indeterminant, and the marginal physical product of $\boldsymbol{\ell}$ is always equal to one. These assumptions, taken together, guarantee that the economy is operating on its PPF. The focus then turns to what can be said about the allocative efficiency characteristics of the point on the PPF where the perfectly competitive economy happens to land.

The economy is assumed to contain a large number, N, of identical Individuals. Each individual owns an endowment, $\boldsymbol{\ell}$**,** of the sole productive resource in the economy. For simplicity, this section assumes that each individual owns one unit of $\boldsymbol{\ell}$; thus, the total amount of the productive resource in the economy, L, is equal to N. The $\boldsymbol{\ell}$'s are sold on the $\boldsymbol{\ell}$-market, where they are purchased by firms that use the $\boldsymbol{\ell}$'s to produce consumption goods, x_1 and x_2, using one $\boldsymbol{\ell}$ for each x_1 or x_2 produced. This production structure implies that the sum of the quantities of x_1 and x_2 will equal L, or equivalently N. Given this simple production structure, not only are the number of firms indeterminant, but firms do not need to specialize in the production of one good or the other. One necessary assumption, however, is that markets are competitive, so the number of firms must be large enough such that all firms are price-takers in all markets.

---

[8] A variant of the model has $U = x_1^{\alpha}(1-\ell)^{(1-\alpha)}$, where $\boldsymbol{\ell}$ is labor and budget constraint $p_1X_1 = w\ell$, in which case $\boldsymbol{\ell}$ will equal α when leisure, $(1-\ell)$, is the numeraire so that w is set equal to one. In this way the labor- leisure choice problem can be introduced into the analysis, if desired.

Consumer Demands

Individuals derive utility from consumption of goods x_1 and x_2 according to a simple Cobb-Douglas utility function:[9]

$$U = x_1^{\alpha} x_2^{1-\alpha}, \;\; 0 < \alpha < 1 \tag{1}$$

The typical individual purchases the goods at market prices p_1 and p_2 using income, w, received from the sale of their single unit of $\boldsymbol{\ell}$. Individuals are price-takers in all markets. The typical individual's budget constraint is therefore:

$$p_1 x_1 + p_2 x_2 \leq w \tag{2}$$

The utility maximization problem generates the following two individual demands:

$$x_1^d = \frac{\alpha w}{p_1} \tag{3}$$

$$x_2^d = \frac{(1-\alpha) w}{p_2} \tag{4}$$

Since there are N (identical in this first case) individuals in the economy, the market demands are simply,

$$X_1^D = \frac{N \alpha w}{p_1} \tag{5}$$

$$X_2^D = \frac{N(1-\alpha) w}{p_2} \tag{6}$$

[9] The following exposition ignores income effects, which is only technically justified when demands are quasi-linear, arising from a quasi-linear utility function [see Hurwicz, 1995]. The Cobb-Douglas form is more familiar and ignoring income effects here does not diminish the generality of concepts that the model is meant to highlight. The same exercise can of course be carried out with quasi-linear utility.

The ℓ market

As mentioned above, the rudiments of production theory can be suppressed by assuming that the ℓ available in the economy is costlessly transformed, one to one, into either of the consumption goods.[10] Individuals sell their ℓ into the ℓ-market where the ℓ's are purchased by firms that transform each unit of ℓ into one or the other of the consumption goods. The ℓ-market is competitive and generates a market clearing equilibrium price, w, for a unit of ℓ. In each consumption-good market, the total amount of good i = 1, 2 produced, X_i^s, is the sum of the outputs of the individual j firms producing good i, which in turn is equal to the total ℓ_i used by the firms producing good i.

$$X_i^s = \Sigma x_{ij}^s = \Sigma \ell_{ij}, \qquad i = 1, 2 \quad j = 1 \ldots . J \tag{8}$$

Where the $\Sigma \ell_{1j} + \Sigma \ell_{2j} = L$.

The Production Possibilities Frontier

If each of the N individuals in the economy owns exactly one ℓ, and each ℓ can be used to produce one unit of good one or one unit of good two[11] the economy's production possibility frontier (PPF) is simply,

$$N = X_1 + X_2 \tag{9}$$

[10] One might also imagine individuals transforming their ℓ endowments into the consumption goods and consuming their own home production. In such a scenario the market for ℓ would not exist. An ℓ-market construct as presented here conveys the notion of markets signaling marginal values of ℓ to market participants and efficiently allocating the ℓ's across production alternatives.

[11] The CRS assumption implies smooth continuous production, so fractional amounts of ℓ can produce equivalent fractional amount of either good.

as depicted in Figure 2.

Figure 2
The Production Possibilities Frontier

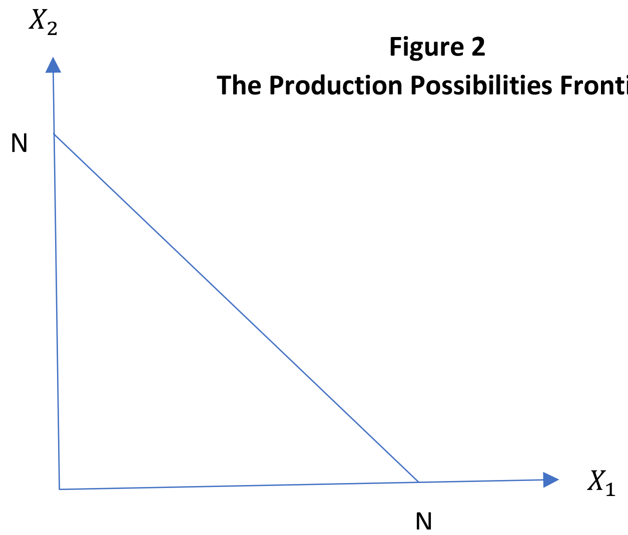

The marginal rate of transformation along the PPF is of course equal to negative one.

The (implicit) Supply Curves

In an economy with perfectly competitive factor and product markets and pure private goods, each point along each a market supply curve reflects the social marginal cost of production. Since the PPF in this model has a slope of negative one, this might seem to imply that the social marginal cost of each good is one. Nevertheless, while it is true that the production of one additional unit of one good always involves the production of one less unit of the other good, this does not imply that the social marginal cost of production of either good is necessarily equal to one. As discussed above, the reason is that the marginal social cost of an incremental unit in one market is derived from the marginal value forgone in the alternative market, which in turn depends on the height of the demand curve in the alternative market and ultimately on marginal utilities. This is the fundamental point driving this exposition.

Each point on a social marginal cost curve reflects the society's social marginal value of resources that would have been available to produce in alternative markets. The social value of those resources in the alternative market is equal to the value of the good or service that would have been produced, and that value is derived from the height of the demand curve in the alternative market. Each social marginal cost curve is therefore derived from the forgone values reflected by the incremental height of the market demand curves in the alternative markets and each social marginal cost curve derived in this way can be thought of as an implicit market supply curve.

This model makes it easy to see that the social marginal cost of an incremental unit in one market is derived from the marginal social valuation of the incremental output reduction in the other market, which is measured by the *demand* value, or willingness to pay, in the other market. The marginal rate of transformation is indeed one along the PPF because production of one more unit of one good always requires producing one less unit of the other, but the marginal social valuation of an incremental unit lost is derived from demand curves, which reflect consumer preferences and marginal utilities. When there exists a lot of one good relative to the other, the marginal social value of the abundant good is low and the marginal social value of the less abundant good is high, reflecting diminishing marginal rates of substitution along indifference curves.

When the economy has a relatively large amount of good two, for example, it is operating relatively far down the good-two demand curve, so the marginal social valuation of a unit reduction of good two will necessarily be lower than when the economy is producing relatively small amounts of good two.

In this way***, the marginal social cost, or the implicit market supply curve, for good one can be derived from the market demand curve for good two; and vice-versa, the implicit market supply curve for good two can be derived from the market demand curve for good one.*** The two marginal social cost curves are each simply geometric reflections of the market demand curves for the alternative goods. The market demand for good two is reflected into p_1-X_1 space to create the implicit market supply curve for good one and the market demand for good one is reflected into p_2-X_2 space to create the implicit market supply curve for good two.

Using (6), the implicit market supply curve for good one, X_1^s, is therefore:

$$X_1^s = N - \frac{(1-\alpha)wN}{p_1} \tag{10}$$

Equivalently, using (5), the implicit market supply curve for good 2, X_2^s, is:

$$X_2^s = N - \frac{\alpha wN}{p_2} \tag{11}$$

Note that p_1 and p_2 are in the denominators of (10) and (11), respectively, because X_1^s exists in p_1-X_1 space and X_2^s exists in p_2-X_2 space. But $p_1$along the supply will be equal to height of the demand for good 2, which is p_2, at each possible allocation ofX_1 and X_2. Deriving p_2 from (6), and substituting it for $p_1$in (10), one gets:

$$X_1^s = N - \frac{(1-\alpha)wN}{\frac{(1-\alpha)wN}{X_2^D}} = N - X_2^D \tag{12}$$

3. Equilibrium

One of the key features of this model as a pedagogical device is that the equilibrium can be numerically derived by specifying the consumer preference parameter, α, and by stipulating the number of individuals in the economy. Since the uniform ***ℓ***'s owned by individuals are the only inputs into the economy and have no intrinsic utility value of their own, there can be no misallocation of inputs in production and all the ***ℓ***'s will all be placed into employment. Thus, as mentioned above, this ensures that the economy is operating on its Production Possibilities Frontier.

Not all points on the PPF are necessarily Pareto optima, of course. With explicit utility functions, spreadsheet software can be used to calculate the utility levels at each point of the PPF. For the Cobb-Douglas case presented here, however, utilities can be aggregated into a single individual, in which case one can simply search for the maximum utility level. For example, when α = 0.5 and N = 100, the utility levels along the PPF are shown in Figure 3. Utility is highest at 50.00, when $X_1 = 50$ and $X_2 = 50$. To find the free-market equilibrium, one need only set (5) equal to (10) and (6) equal to (11), choose w as the numeraire, and solve to find that the free-market equilibrium results in $P_1 = P_2 = w = 1$ and lands at point $X_1 = X_2 = 50$, the point where social utility is indeed maximized.

Figure 3

Utility Level at each integer point on the PPF		
Utility	PPF	
	X1	X2
0	0	100
9.95	1	99
14.00	2	98
.	.	.
.	.	.
49.99	49	51
50.00	50	50
.	.	.
.	.	.
9.95	99	1
0.00	100	0
0.00	100	0

Of course, it is the inspection of total utility at an out-of-equilibrium point and the numerical calculations for movement along the PPF toward the perfect competition equilibrium that provides the greatest revelations to the student. Consider, for example, one such out-of-equilibrium point for an economy in which N = 100 and α = .5. The two markets are depicted in Figure 4. The society has one unit of X_1 and ninety-nine units of X_2. The point at the price $P_1 = 50/99$ on the supply curve in Market One is a reflection of the "50/99" valuation point on the demand curve in market two. With w as the numeraire, the marginal social cost of an incremental unit of good two is 50, and the marginal social cost of good one, derived from (11), is 50/99.

The PPF is still a straight line with endpoints at 100 on each axis and the production tradeoff is indeed one to one, but when the economy is producing ninety-nine units of good two and only one unit of good one, ***the marginal social cost of an incremental unit of good two is the very high forgone value***

of single unit of good one, and it is derived from (10), which itself comes from (6). As the amount of good two in the economy falls, its marginal social valuation rises along its demand curve, causing a rise in the marginal social cost of good one, thus implying the upward sloping marginal social cost (supply) curve for good one, represented by (10).

With identical households, total utility in this economy at this point on the PPF ($X_1 = 1\ , X_2 = 99$) is as listed in Figure 3,

$$U = 1^{.5}99^{.5} \approx 9.95 \tag{13}$$

Figure 4

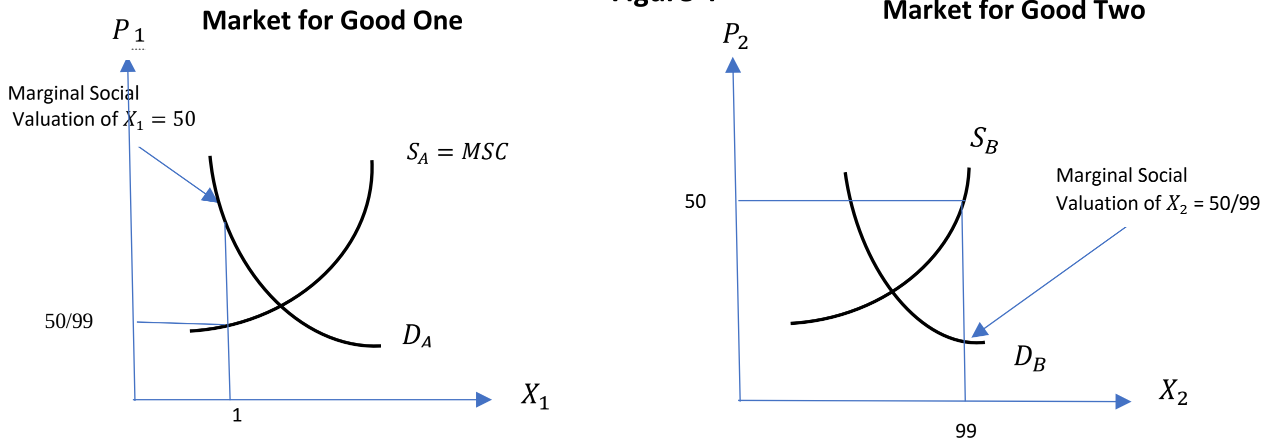

Of course, the point $X_1 = 1\ , X_2 = 99$ is on the PPF but is not the perfect-competition market equilibrium. If price $P_1 = 50/99$ were to be imposed in Market One, there would be a large excess demand in Market One, as can be seen in Figure 4, and if a price of 50 were to prevail in Market Two, there would be a large excess supply in Market Two.

Ultimately, this numerical illustration demonstrates production efficiency of the perfect-competition equilibrium and highlights that the market equilibrium point on the PPF is a Pareto optimum, the first fundamental theorem of welfare economics. In addition, in the current example, with N identical consumer and an aggregated Cobb-Douglas utility, it is also the utility maximum or social-welfare maximizing point.

A simple change in the preference parameter will prompt a new calculation and will naturally lead to a different equilibrium point. When α = .2, for example, the market equilibrium will pick out point $X_1 = 20$ and $X_2 = 80$, on the PPF. Thus, the model provides the instructor the opportunity to change parameters and provide numerical exercises to students to help deepen their understanding of these fundamental economic concepts.

4. Inequality and the second fundamental theorem

Consider next, income or wealth inequality as an alternative area of enquiry. This time assume that N = 80, but with each memberof the *wealthy* twenty-five percent of the population owning twice the $\boldsymbol{\ell}$-endowment as a member from the less wealthy 75 percent. In other words, twenty individuals each own two $\boldsymbol{\ell}$'s, while the other sixty individuals each own only one $\boldsymbol{\ell}$. Thus, there continue to exist 100 units of $\boldsymbol{\ell}$ in the economy, so the PPF in this world has not changed and still has end points at L = 100 on each axis (although L is no longer equal to N), and the slope of the PPF is still equal to one. Assume, for the moment that individuals have identical preferences, say at α = 0.5 as in section 3. Then the general structure of demand is unchanged. Each wealthy individual has double the demand of a typical less wealthy individual. The twenty wealthy have the demand of forty individuals from section three above and the sixty less wealthy individuals have the same demand.[12]

In other words, the market demands are identical to the market demands of section 3, and the equilibrium lies again at $Q_1 = Q_2 = 50$. The difference is that each member of the wealthy twenty-

[12] The wealthy, having two $\ell's$, will have demands as in (3) and (4) but multiplied by two, while the less wealthy will still have demands exactly as in (3) and (4).

five percent of the population will receive one unit of each good, while each member of the less wealthy seventy-five percent will receive one half unit of each good, as in section 3.

This example demonstrates that unequal initial endowments or inequality in general is not addressed by the Pareto optimality concept. Prices allocate resources efficiently and a Pareto optimum is achieved, but this result should lead to questioning whether the market equilibrium is a *social optimum* and whether society's social welfare function might address the inequality issue and move society away from the particular Pareto optimum picked out by the competitive equilibrium of the market. This can lead to a discussion about whether a society in which taxes and redistributions create production inefficiencies, leaving society inside its PPF, might nevertheless be at a more socially desirable point than at the original market-chosen Pareto optimum point on the PPF.

Taking this line of questioning a step further, a class might consider what would happen if the wealthy and less wealthy have different preferences, say α = 0.2 for the wealthy twenty individuals and α = 0.8 for the sixty less wealthy individuals. Under this new set of preferences, it is easy to show that the wealthy individuals will each consume 0.4 units of good one and 1.6 units of good two, while each of the less wealthy individuals will consume 0.8 units of good one and 0.2 units of good two. The demand and supply schedules for this scenario can again be programmed into a spreadsheet to show that this perfectly competitive economy will produce 56 units of good one and 44 units of good two. In other words, the purchasing power of the wealthy causes production and society's resources to shift in the direction of the preferences of the wealthy and away from the production of those goods most strongly preferred by the less wealthy. Reallocating wealth across society, by changing the ℓ endowments, will move the economy to various Pareto optimal points along the PPF, thus illustrating the second fundamental theorem of welfare economics – any Pareto optimum can be achieved as a competitive equilibrium after appropriate rearrangements of the original endowments.

5. Conclusion

The simple computational general equilibrium model presented above is a new pedagogical device. As of this writing, the author has not seen it used in modern microeconomic textbooks or in lectures or PowerPoint presentations. The above presentation suggests that this model can be an important and enlightening complement to the traditional Edgeworth-Bowley Box illustrations of the tangency conditions of the fundamental theorems of welfare economics. The model is analytical but also allows for numerical simulations. In a two-good world, the use of the reflection of a market demand curve as the marginal social cost, or implicit supply curve, in the alternative market, allows one to bypass much of production theory and directly focus on the first and second theorems of welfare economics, the Pareto efficiency of the competitive equilibrium, and make the important point that the Pareto efficient competitive equilibrium might not be the social optimum.

The results presented above show that illustrative numerical examples can be used to provide a more concrete presentation of the general concepts and can also be used to provided practical exercises for students to work through. The model complements the Edgeworth-Bowley Box presentations by providing market equilibria and Pareto optimal solutions numerically, allowing for simple simulations for a variety of scenarios, including inequality of initial endowments and alternative preference configurations. In the hands of enthusiastic educators, it can be modified in any number of ways to highlight the nuanced issues of societal well-being and welfare economics.

References

Arrow, Kenneth J.; Debreu, Gerard, *"Existence of an Equilibrium for a Competitive Economy"*. *Econometrica*, ***22*** *(3)* July 1954: *265-290.*

Bowley, Arthur W. *The Mathematical Groundwork of Economics*. Oxford: Clarendon Press, 1924, reprinted, New York: Augustus M. Kelley, 1965.

Edgeworth, Francis Y. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. London: Kegan Paul, 1881. http://onlinebooks.library.upenn.edu/webbin/book/browse?type=lccn&key=HB74%20%2eM3%20E3%201881 , accessed January 11, 2021.

Hicks, J. R., The Foundations of Welfare Economics, *The Economic Journal*, Vol. 49, No. 196 Dec. 1939, pp. 696-712

Hotelling, Harold, "The general welfare in relation to problems of taxation and of railway and utility rates," Econometrica, July 1938.

Humphrey, Thomas M., The Early History of the Box Diagram, Economic Quarterly, Volume 82/1, Federal Reserve Bank of Richmond, Winter 1996, pp37-75.

Hurwicz, Leonid, What is the Coase theorem? *Japan and the World Economy*, 7:60–74, 1995.

McKenzie, Lionel W., "On Equilibrium in Graham's Model of World Trade and Other Competitive Systems", *Econometrica*, 1954.

Stiglitz, Joseph E., (March 1991), "The Invisible Hand and Modern Welfare Economics", NBER Working Paper No. 3641,