

Why Christians should Repudiate Rational Egoist Capitalism: Perfect Competition without morality cannot guarantee efficiency

I. A Proof.

Imagine an island society of purely selfish agents. In this society adult agents either hunt, gather, or farm during a 6-day period. You might say, agents are self-sufficient “Robinson Crusoes” who produce and consume for themselves. While each of the three types of commodities associated with each activity are available throughout the region in which the population dwells, there exists a variation in distribution in each type. The result is that there also exists a variation in the amount of effort it takes to acquire a given commodity. On the seventh day, agents gather to trade some of their produce for other commodities they desire. Prices in this society are simple exchange ratios. In other words, the price of a unit of one commodity is given in terms of a number of units of another commodity. Thus, in order to acquire a certain amount of one type of commodity agents give up or trade set amounts of other commodities. The price of a loaf of bread, for example, may be one fish. No agent can affect prices except through the conventions of a price mechanism. In this case, agents take prices as they are given by a *Walrasian auctioneer*¹ who, in the presence of all prospective traders, calls out a set of prices. Each trader then, on the basis of his preferences, determines what amount of each commodity he or she is willing to buy or sell at those prices and reports this to the auctioneer. Unless all commodities which agents desire to sell are sold, and all which agents desire to buy are acquired, the auctioneer calls out a revised set of prices - lowering (raising) the exchange ratio on those unsold commodities (commodities not offered for sale).

For the sake of rigor and proof, let us call this social situation: “Strict Rational Egoism”.² Here are the crucial features:

- (p₁) *Agents preferences range over alternative social states defined solely in terms of consumption bundles.*
- (p₂) *Agents’ preference relations are stable, rational, and locally non-satiated.*
- (p₃) *Agents’ goals are selected according to a utility maximization criterion.*
- (p₄) *Agents’ beliefs depend only on information.*
- (p₅) *Agents are sufficiently and instrumentally rational.*

¹ The *auctioneer* is an element in Leon Walras’ (1874) seminal analysis of competitive general equilibrium.

² Background and expansion of this argument may be found in the author’s book, *The Moral Conditions of Economic Efficiency*, (Cambridge University Press, 2001).

(p₆) *Agents are constrained by a perfectly competitive market: numerous participants, homogeneous products, freedom of exit and entry, and perfect information.*

(p₇) *Agents control finite resources.*

(p₈) *There are no moral rules.*

(p₉) *There are conventions to equilibrate supply and demand.*

Some people believe that if agents possess maximal information-processing capabilities and perfect information regarding *every economically-relevant variable*, non-market actions (i.e., force and fraud) will be precluded and equilibrium allocations will be Pareto optimal because every agent will know that everyone is able to anticipate and to neutralize non-competitive behavior and every agent also knows that everyone else knows this. If so, then in theory, moral normative constraints will not be necessary. All that is needed is better information. However, if it can be shown that, in social situations defined by assumptions (p₁) - (p₉) there exists the possibility that agents will not be able to decide what to do, then perfect information is not sufficient to secure competitive. Moreover, it will also follow that no information-revealing mechanism for strict rational egoism can rectify the result. Only a set of moral normative constraints will make Pareto optimal equilibrium allocations possible for strict rational egoists - even in the theoretical ideal.³

Perfect information and full rationality

By the term, “perfect information”, I mean complete and symmetrically-distributed information regarding every economically-relevant variable. That is to say, each individual is aware of the complete social situation in which he finds himself; no factor of the situation which could affect any individual’s utility level is hidden. Stated in the formal terms we have adopted, we say that each agent knows each economically-relevant variable pertinent to every agent including every element of the entire *social situation*.⁴ Because perfect information covers every economically-relevant variable, it entails common knowledge. That is, each agent knows *that every other agent knows* the complete social situation including the fact that everyone knows that everyone knows, and so on. To gather all of this together in one concise statement of three broad categories of knowledge, we say that each agent knows (1) every relevant characteristic regarding every *agent* including, of course, himself,

³ This clearly shows the essential role of moral normative constraints and the detrimental consequences of not rigorously specifying and modeling them in the First Welfare Theorem.

⁴ Our account follows Ostrom *et al.* (1994:34). It is also compatible with the concept of a “social situation” in Greenberg (1990).

(2) the value of every relevant variable in the *situation*, (3) that every other individual knows (1) and (2), and (4) that everyone knows (3). In what follows I develop these ideas in stages and formalize them in order to provide a rigorous and clear statement of the problem.

We can represent the information-processing capability of each agent as an abstract computer which “remembers” sets and relations, and which is programmed to enumerate any effectively enumerable set or relation and to compute any effectively computable function. Informally, we say that a function is *effectively computable* if and only if there is an algorithm such that when given any element from the domain of the function as input, the algorithm gives as its output the unique element from the range of the function. Accordingly, we approach the central question before us using some concepts and established results in computability theory. Two established results are crucial to this question. First of all, let us suppose that individuals’ information processing capabilities are Universal Turing Machines. A *Universal Turing Machine* is an abstract computer in which the limitations attached to actual computers (e.g., time, speed, and material) are irrelevant, and which can compute every *effectively computable* function.⁵ Imagine it this way: in every respect, individuals’ cognitive capacities are those of a normal human being, except that each can store information and compute functions better than normal human beings. But individuals in our model differ from computers and are similar to humans in that each is capable of second-order awareness; they are aware *that* they are perceiving, computing, desiring, and the like. A second standard result in computability theory which is integral to our argument is that the class of effectively computable functions is coextensive with the class of Turing-computable functions.⁶ Therefore, for any function f , individuals in our model can (cannot) compute f , if f is (is not) effectively computable.⁷

The particular course of action an individual takes, therefore, is a matter of that individual’s computing a function. Let a *decision function* be the function that each individual

⁵ An easily accessible explication of the concept of a *Universal Turing Machine* is given in Beckman (1980:183,4). Alternative proofs are given in Davis (1973:64), Malitz (1979:95-9), and Yasuhara (1971:28-30).

⁶ This is Alan Turing’s (1936) thesis. See also Boolos and Jeffrey, (1989:20,54).

⁷ I need not prove that each function is effectively computable and each set or relation is effectively enumerable. Nevertheless, perhaps this much is worth noting: proving that a function is effectively computable or that a set or relation is effectively enumerable using Turing Machine descriptions is cumbersome. Since the class of partial recursive functions is coextensive with the class of partial Turing-computable functions, it would be less cumbersome to show that, for any function f , individuals in my model can (cannot) compute f , if f is (is not) partial recursive.

computes in order to be able to decide what action to take. The elements of each individual's decision function is specified, given the factors of information individuals have in their possession. I will show that there exists a set of decision functions which are not effectively computable. Thus, even supposing perfect information and flawless information processing purely selfish agents are not guaranteed to achieve the common good.

Information Units

We have before us four broad categories regarding which individuals possess perfect information. Each agent knows (1) every relevant characteristic regarding every *agent* including, of course, himself, (2) the value of every relevant variable in the *situation*, (3) that every other individual knows (1) and (2), and (4) that everyone knows (3). In order to describe adequately the members of these categories and how they affect individuals' decisions, we must relate these categories of individuals' knowledge to the analytical framework within which we are working and to the premises of the First Welfare Theorem. Recall that we are now treating assumptions (p₁) - (p₉) as the assumptions of the version of the First Welfare Theorem. Since every agent possesses perfect information each individual is fully aware of each assumptions (p₁) - (p₉) - not *as* an element in an analytical model of social interaction nor *as* a premise in a proof, but as a feature of the situation in which they act.

Viewed as features of the situation in which they act, every element of the *agent* category or the *situation* category varies in its instantiation with respect to each individual, thereby contributing to the difference that exists between different individuals' decision-pertinent data. For example, premise (p₇) is *Agents control finite resources*. But from the perspective of each individual, even though each individual faces the *fact* that each control finite resources, each individual set of resources varies in its content. There are seven such elements or factors from the *agent* category and the *situation* category which are directly pertinent to each individual's decision.

Factor 1. Everyone knows that every individual is fully rational in that each is maximally capable of storing and processing information.

Factor 2. Everyone knows the set of all individuals in the action arena.

Factor 3. Everyone knows the set of strategies feasible for every individual.

Factor 4. Everyone knows the set of combinations of all strategies.

Factor 5. Each individual knows the set of potential outcomes achieved by alternative combinations of strategies.

Factor 6. Every agent knows every individual's utility function defined on the set of potential outcomes.

Factor 7. Every individual knows factors (1) - (6), and that everyone knows factor (7), that is, that everyone knows that everyone knows.

We can then provide detail to each of these factors in order to picture more easily the decision process of strict rational egoists.

Factor 1: Everyone knows that every individual is fully rational in that each is maximally capable of storing and processing information. Therefore, each individual must take into consideration and calculate every other individual's decision function because the situation is strategic and non-cooperative.

Factor 2: Every individual knows the set of all individuals in the action arena. This means that every individual has a list containing the names of all individuals, can determine how many individuals are on the list, and can pick out any individual from the list. Let the set of individual's names be denoted by $\mathbf{I} = \{1, \dots, n\}$. Each individual i has a complete list of each $i \in \mathbf{I}$ in his memory, as follows: each $i \in \mathbf{I}$ is assigned exactly one positive integer beginning with the positive integer, 1. No two individuals are assigned the same integer. Thus, $\mathbf{I} = \{1, \dots, n\} \subseteq \mathbb{N}$. Let $\max(\mathbf{I})$ be the effectively computable function that gives the maximum value of \mathbf{I} . Let $f_i : \mathbb{N} \rightarrow \mathbb{N}$ be the effectively computable function which associates exactly one positive integer with a name of an individual; no two individuals have the same name.⁸

Thus, every individual knows how many agents are in \mathbf{I} because he can compute $\max(\mathbf{I})$, and every individual i can identify (i.e., pick out) any $i \in \mathbf{I}$, $i = 1, \dots, n$, because he can compute f_i .

Factor 3: Every individual i knows the set of strategies feasible for each $i \in \mathbf{I}$, which we call i 's *natural strategy domain*, and denote as S^i . This means that, for each S^i , every individual has a list containing the index numbers of each strategy in S^i , knows how many strategies are on each list, can pick out any strategy from any list. For the sake of simplicity and without loss of generality, we define a *strategy* a_n^i as the n -th single course of action available to

⁸ Technically, speaking the domain of the function $f_i : \mathbb{N} \rightarrow \mathbb{N}$ is a finite subset of *names of positive integers*. Hence, for each individual in $\mathbf{I} = \{1, \dots, n\} \subseteq \mathbb{N}$, the function associates a *name* of a positive integer with a name of an individual; no two individuals have the same name. But we will ignore this technicality.

individual i under a particular set of conditions.⁹ Each individual has a complete list of each $a_n^i \in S^i$ in his memory, as follows: each $a_n^i \in S^i$ is assigned exactly one positive integer beginning with, a_n^i , and no two strategies are assigned the same integer. Hence, for each individual i , $S^i = \{a_n^i, \dots, a_n^i\}$, where each n is its index number. Let $\max(S^i)$ be the effectively computable function that gives the maximum index number in S^i . Let $f_S^i : k \rightarrow S^i$, for $k = \{1, \dots, n\}$, be the effectively computable function that associates a single positive integer with each $a_n^i \in S^i$.

Thus, every individual knows how many strategies are in S^i because he can compute $\max(S^i)$, and each individual can identify any $a_n^i \in S^i$ because he can compute f_S^i . Each individual knows the class $K = \{S^1, \dots, S^n\}$ of all natural strategy domains S^i , $i = 1, \dots, n$. Each individual holds in memory a complete list of each $S^n \in K$, such that exactly one positive integer is assigned to each strategy domain beginning with, S^1 , until the list, S^1, \dots, S^n is complete, and no two natural strategy domains are assigned the same integer. Furthermore, each individual knows how many natural strategy domains S^i are on the list, and can pick out any strategy domain from the list K or any strategy from any natural strategy domain. Let the effectively computable function $f_K : k \rightarrow K$, for $k = \{1, \dots, n\}$, be the function that associates a single positive integer with each $S^i \in K$. Therefore, each individual can identify any $S^i \in K$, because he can compute f_K , and each knows how many are on the list because each knows how many individuals there are and their "names".

Factor 4: Everyone knows the set of combinations of all strategies. Let $SX = \{s_1, \dots, s_n\}$ be the Cartesian product of the S^i 's, where each s_n represents one *strategy n-tuple* in SX . Let $h: SX \rightarrow Z$ represent the effectively computable outcome function, where Z is the outcome space. In other words, h maps each total combination of actions, $s_n \in SX$, into the outcome space Z . Each individual knows the unique social state achieved by any $s_n \in SX$ only if he knows SX and he can compute h . This means that every individual has a list containing the index number of all $s_n \in SX$, knows how many social states are on the list, and can pick out any social state from the list. Each individual has a complete list of each $s_n \in SX$ in his memory, as follows: each $s_n \in SX$ is assigned exactly one positive integer beginning with, s_1 , and no two collective strategies are assigned the same integer. SX is effectively enumerable. Let the

⁹ The "conditions" to which we refer can be conceived of as type-variables, so that individuals are not required to know every conceivable configuration of conditions, but only the features of conditions that affect his choice of a strategy. Sets of actual conditions are members of classes (types) of conditions differentiated by their respective relevant defining features.

effectively computable function $f_{SX} : k \rightarrow K$, for $k = \{1, \dots, n\}$, be the function that associates a single positive integer with each $s_n \in SX$. Let the effectively computable function $h : SX \rightarrow Z$ map the set of strategy n -tuples onto the set of outcomes. The function h is a 1 - 1 onto function. The function h is a 1 - 1 because $\forall s_n, s_m \in \text{dom}(h), s_n \neq s_m \Rightarrow h(s_n) \neq h(s_m)$, and onto because $\text{dom}(h) = SX$, and $\text{ran}(h) = Z$.

Each individual can compute the Cartesian product of two sets, given correct information regarding the elements of each set. In this case each individual knows how many individuals there are and how many actions are available to each. Given this fact, we observe that each individual can pick out any strategy n -tuple because he can compute f_{SX} . Furthermore, because each individual can compute the function h for each action $s_n \in SX$, every agent knows the social state σ_n for which s_n is necessary and sufficient. But he must be able to determine $a_n^i \in s_n$, which indicates his own part in achieving the σ_n associated with s_n . Let $g_n^i : s_n \rightarrow a_n^i$ represent the effectively computable function which picks out i 's action in the collective strategy s_n . There are, of course, several such functions; one for each s_n .

Factor 5: Each individual knows the set of *potential outcomes*, referred to alternatively as *social states*. This means that every individual has a list containing the names of all social states, knows how many names of social states are on the list, and can pick out any social state from the list. Social states are defined by alternative allocations of commodities. Let the set of names of possible social states be denoted by $Z = \{\sigma_1, \dots, \sigma_n\}$, where each n is its index number. Each individual has a complete list of each $\sigma_n \in Z$ in his memory, as follows: each $\sigma_n \in Z$ is assigned exactly one positive integer beginning with, σ_1 , and no two social states are assigned the same integer. Let $\text{max}(Z)$ be the effectively computable function that gives the maximum value of Z . Let $f_Z : k \rightarrow Z$, for $k = \{1, \dots, n\}$, be the effectively computable function that associates a single positive integer with each $\sigma_n \in Z$.

Thus, every individual knows how many social states are in Z because he can compute $\text{max}(Z)$, and each individual can identify any $\sigma_n \in Z$ because he can compute f_Z .

Factor 6: Every agent knows every individual's utility function on Z . For the sake of clarity of presentation, we first show that each individual knows his or her own utility function. Technically speaking, a *utility function* $u(x)$ assigns a numerical value to each member of a set X of alternatives. For our purposes, X is a set of alternative social states.

We begin with the supposition that each agent knows his or her own preference relation on Z . Notice that Z is not a set of alternative consumption bundles, but rather is a set of alternative social states characterized in part by alternative consumption bundles.

Stipulating that the symbol, “ \succeq_i^Z ”, represents individual i 's preference relation on Z , we assume that \succeq_i^Z is reflexive, transitive, complete, and continuous as we did in chapter two.

Every preference relation having these four properties can be represented as a utility function.¹⁰ Let $U_i: Z \rightarrow B \subseteq \mathbb{R}$, for all i , where \mathbb{R} is the set of real numbers, and B is a proper subset of \mathbb{R} , represent an effectively computable individual utility function. We want to indicate that for any two social states, individual i thinks that the first is as least as good as the second if and only if the utility value of the first is greater than or equal to the second. Hence, each $U_i: Z \rightarrow B \subseteq \mathbb{R}$ is a utility function such that

$$\forall \sigma_m, \sigma_n \in Z [\sigma_m \succeq_i \sigma_n \Leftrightarrow U_i(\sigma_m) \geq U_i(\sigma_n)].$$

Let u_m^i be a real number representing the value i 's utility function U_i , for some argument, σ_m , that is, $U_i(\sigma_m) = u_m^i$. Each $u_m^i \in \text{ran}(U_i)$, that is, each $u_m^i \in B \subseteq \mathbb{R}$ can be given an ordinal ranking, beginning with the integer, 1. Let the effectively computable function $f_B: B \rightarrow A \subseteq \mathbb{N}$, be a 1-1 onto function whose domain = $\text{ran}(U_i)$ and whose range = $\{1, \dots, n\}$, such that

$$\forall u_m^i, u_n^i \in B, \forall n, m \in A [u_m^i > u_n^i \Leftrightarrow n > m].$$

Hence, $f_B(u_m^i) = n$. Let the function $f_A: A \subseteq \mathbb{N} \rightarrow Z$, be a 1-1 onto function whose domain = $\text{ran}(f_B)$, and whose range = Z . Hence, $f_A(n) = \sigma_n$.

Note that each function U_i is a 1-1 onto function. U_i is 1-1 because $\forall x, y \in \text{dom}(U_i)$, $x \neq y \Rightarrow U_i(x) \neq U_i(y)$, and U_i is onto because $\text{dom}(U_i) = Z$, and $\text{ran}(U_i) = B$.

For each outcome $\sigma_n \in Z$, each individual i knows its ordinal rank of n possible alternatives for each $i \in I$, because i can compute each U_i, f_B , and f_A .

Factor 7: Every individual knows factors (1) - (6), and that everyone knows factor (7), that is, that everyone knows that everyone knows. At the risk of redundancy, we state the same idea in other words: each agents knows (1) that everyone knows that every individual is a utility-maximizing, instrumentally rational, Turing-machine calculator, (2) that every individual has a list containing the names of all individuals, knows how many individuals are on the list, and can pick out any individual from the list, (3) that every individual has a list of containing the index number of all social states, knows how many social states are on

¹⁰ See Varian (1992:95-7).

the list, and can pick out any social state from the list, (4) that everyone knows everyone's utility function, (5) that everyone knows everyone's feasible actions, (6) that everyone knows the outcomes that follow from those actions, and (7) that everyone knows that everyone knows (1) - (6).

Parametric choice

For each strategy n -tuple, $s_n \in SX$, there exists a social state σ_n for which s_n is both necessary and sufficient, while $a_n^i \in s_n$ is only necessary. But where, for some individual i , the actions of agents other than i are irrelevant, a_n^i is both necessary and sufficient to achieve σ_n . Let a *parametric decision function* be a type of decision function in which common knowledge is not a factor and in which the actions of other agents are irrelevant.¹¹ Hence, where other individuals' actions are irrelevant, an agent maximizes utility if and only if he takes that action a_n^i which achieves that social state holding the highest level of utility. In other words, the particular decision that an individual makes (i.e., the strategy that he pursues) is the one that maximizes utility. Therefore, in order to decide what action to take, each individual must make the following determinations. First of all, an individual must determine what social state gives the highest level of utility, u^i . Recall that each $U_i: Z \rightarrow B \subseteq \mathbb{R}$. Each agent thus, must ascertain the highest value in B . Let the effectively computable function, $\max(B_i)$ pick out the highest value in B .

By computing $\max(B_i)$, f_B , and f_A , each agent can determine what social state gives him the highest level of utility. Secondly, each individual must determine what collective strategy achieves that outcome by computing h^{-1} . Finally, each agent must determine what individual action is a member of that collective strategy by computing g_n^i . Therefore, we represent i 's decision as the value of his parametric decision function as follows:

$$a_n^i = g_n^i(h^{-1}(f_A(f_B(\max(B_i))))).$$

Proposition: Every parametric decision function $g_n^i(h^{-1}(f_A(f_B(\max(B_i)))))$ is effectively computable.

Notice that the first step in a decision process is finding the social state with the highest ordinal value. Perhaps an explanation for this claim is in order. Some types of relevant data (some of which are given and others of which require calculation or experimentation) are not directly relevant to the argument of this chapter. To depict how agents acquire these types

¹¹ Notice that, by itself, the absence of common knowledge does not render the actions of other individuals irrelevant.

of information would needlessly complicate things. For examples, a list of all individuals and a list of strategies for each is simply given. Only *that* each agent knows the set of all agents and all of their actions is relevant to our argument. Calculating the Cartesian product of all strategy domains is not relevant to the proof. Neither is the manner in which the set of possible outcomes is determined. We only require a mapping of strategy n -tuples onto outcomes having alternative payoffs. Therefore, we simply assume a finite set of outcomes. Finally, each preference relation on the set of outcomes is given and calculating each utility function from its corresponding preference relation is not relevant. With this much information at their disposal, the *next* step individuals take is to find the outcome yielding the highest utility. This step is the first step in the process of computing a decision function.

Quasi-parametric decision. Let a *quasi-parametric decision* be one in which some, but not all, of every other agents' actions in a strategy n -tuple are irrelevant. We will have occasion to make reference to such types of decisions later. We bring it up now because of its relation to both parametric and strategic choice.

Conditional decision. Let a *conditional decision* be a *stage* in a decision process in which an individual must consider a range of hypothetical alternatives before making a final decision. Each hypothetical alternative is a conditional decision. Each agent reasons in this manner: *Suppose I take action a_n^i , and j takes action a_n^j , and so on, social state σ_n will follow.* The kinds of parametric decisions with which we are concerned do not involve conditional decisions in this sense, simply because alternative means are incorporated into the set of social states, so that by choosing the social state having the highest ordinal value, an individual chooses both a goal and the best means to achieve it.

Strategic choice.

For each individual i , the action i takes depends on the maximum value of his utility function, which itself depends in part on the actions taken by every other agent to achieve the maximum value of their respective utility functions. Hence, each individual must take others' decision functions into account. Therefore, strategic decisions involve conditional decisions because individuals first compute a parametric decision function, then determine which *other* strategy n -tuples of which the action a_n^i is a member. Each strategy n -tuple, except one, results in a social state different than the one which gives the calculating agent her highest utility. We give a summary version of each individual's decision function as follows:

$$a_n^i = g_n^i(s_m : g_n^i(h^{-1}(f_A(f_B(\max(B_i)))))) = g_n^i(h^{-1}(f_A(f_B(\max(B_j))))), \forall j \neq i.$$

Decision functions are effectively computable if and only if there exists a Nash equilibrium of strategies, that is, given every other players' strategy, no other strategy yields a higher payoff. Thus, the effective computability of decision functions depends on the character of each utility function. Our concern now is to show that there exists a set of decision functions which are not effectively computable.

There exists a set of utility functions which render these decision functions not effectively computable under conditions of perfect information.

An alternative way to state the proposition is that there is a set of utility functions for which there is no Nash equilibrium of strategies, so that no individual's decision function is effectively computable. Suppose, the *Walrasian auctioneer* presents each individual i with an initial price vector \mathbf{p}^* . Each individual's natural strategy domain includes three possible alternative actions: defraud by overstatement, trade, or defraud by understatement. That is, for each individual i , $S^i = \{a_1^i, a_2^i, a_3^i\}$, where $a_1^i = +\text{fraud}$, $a_2^i = \text{trade}$, and $a_3^i = -\text{fraud}$. Each individual then computes a decision function to give him the best course of action. Finally, each agent simultaneously presents his trading decision to the *Walrasian auctioneer*, where a trading decision is a list of the amount of each commodity the agent presents himself as willing to trade, given \mathbf{p}^* . We present the case for two individuals which can be extended to n -individuals without loss of generality.

Now, before I present the specifics of the proof and in order to make the proof itself more perspicuous, I will present and respond to a possible objection that perfect information seems to rule out the possibility of fraud. The objection could be argued as follows:

The proof requires that agents' natural strategy domains include two means to defraud. But since the perfect information assumption has been strengthened so as to include common knowledge of every agent's utility function defined over possible outcomes, the possibility of fraud by overstatement or understatement of one's preferences to the Walrasian auctioneer is ruled out. That is, to defraud by either of these means involves stating one's preferences to be what everyone knows them not to be. Thus, fraud is not an alternative to trade as the proof seems to require.

We must keep in mind the difference between *possible actions* which are members of an agent's *natural* strategy domain and *effective actions* which are members of agent's *rational* strategy domain, which is a subset of that agent's natural strategy domain. The fact that each agent knows every other agents' utility function does not render the actions impossible. It only renders them ineffective. Fraud by overstatement or understatement of one's

preferences to the *Walrasian auctioneer* remain possible actions and therefore members of agents' natural strategy domain. Perfect information is supposed to render such possible actions ineffective, thus ruling out non-price taking behavior and leaving trade as the only option. Indeed, in principle, for some or even most sets of utility functions defined over outcomes, perfect information in this framework will eliminate non-price taking behavior thus rendering it sufficient for efficient outcomes without moral normative constraints. However, as the forthcoming proof shows, for this given set of utility functions, even though agents know each other's utility functions, they cannot come to a decision regarding which action to take. Part of each agent's decision process involves first identifying the most preferred outcome and the action which is suppose to achieve that outcome. Then, realizing that everyone else knows his most preferred outcome and its means and that they can respond in such a way as to take advantage of such a move, the agent considers an alternative. But others know this as well. Each agent knows this about every other agent. They cannot, given this particular set of decision functions, achieve that efficient outcome which results when both decide to trade. Thus, perfect information regarding every economically relevant variable and perfect information processing capabilities are not sufficient to achieve economic efficiency. However, the addition of moral normative constraints to that same social situation ensures economic efficiency by ruling out fraud.

Consider now the specifics of the proof. In *Table 1* I list the natural strategy domains, $S^i = \{a_1^i, a_2^i, a_3^i\}$ for each individual i , and the payoff each individual receives for each action a_n^i .

	a_1^2 +fraud	a_2^2 trade	a_3^2 -fraud
a_1^1 +fraud	16,8	9,2	4,14
a_2^1 trade	13,10	6,12	10,6
a_3^1 -fraud	7,18	8,1	12,3

Table 1

Therefore, some of the information that each agent possesses is given here:

1. $I = \{1,2\}$;

2. $Z = \{\sigma_1, \dots, \sigma_9\}$;
3. $U_1 = \{\langle \sigma_1, 9 \rangle, \langle \sigma_2, 5 \rangle, \langle \sigma_3, 1 \rangle, \langle \sigma_4, 8 \rangle, \langle \sigma_5, 2 \rangle, \langle \sigma_6, 6 \rangle, \langle \sigma_7, 3 \rangle, \langle \sigma_8, 4 \rangle, \langle \sigma_9, 7 \rangle\}$,
and
 $U_2 = \{\langle \sigma_1, 5 \rangle, \langle \sigma_2, 2 \rangle, \langle \sigma_3, 8 \rangle, \langle \sigma_4, 6 \rangle, \langle \sigma_5, 7 \rangle, \langle \sigma_6, 4 \rangle, \langle \sigma_7, 9 \rangle, \langle \sigma_8, 1 \rangle, \langle \sigma_9, 3 \rangle\}$;
4. $K = \{S^1, S^2\}$, where $S^1 = \{a_1^1, a_2^1, a_3^1\}$, and $S^2 = \{a_1^2, a_2^2, a_3^2\}$;
5. $SX = \{\langle a_1^1, a_1^2 \rangle, \langle a_1^1, a_2^2 \rangle, \langle a_1^1, a_3^2 \rangle, \langle a_2^1, a_1^2 \rangle, \langle a_2^1, a_2^2 \rangle, \langle a_2^1, a_3^2 \rangle, \langle a_3^1, a_1^2 \rangle, \langle a_3^1, a_2^2 \rangle, \langle a_3^1, a_3^2 \rangle\}$,
where $s_1 = \langle a_1^1, a_1^2 \rangle, s_2 = \langle a_1^1, a_2^2 \rangle, \dots, s_9 = \langle a_3^1, a_3^2 \rangle$;
6. $h = \{\langle s_1, \sigma_1 \rangle, \dots, \langle s_9, \sigma_9 \rangle\}$.

Each individual's strategic decision process

(In what follows, social states named beginning with σ_1 at the top left and moving to the right.) In general, the first step in the process of computing a decision function is to find the outcome yielding the highest utility. Each agent then searches for the strategy n -tuple which achieves that social state yielding the highest level of utility for i and determines whether the strategy n -tuple which achieves it is a Nash equilibrium. According to Factor 5 and Factor 6, each individual knows the set of potential outcomes achieved by alternative combinations of strategies, and every agent knows every individual's utility function defined on the set of potential outcomes. Therefore, each individual begins by making a conditional decision which is aimed at determining what action yields the social state giving the highest level of utility and which, in turn, is achieved by computing the parametric decision function $a_n^i = g_n^i(h^{-1}(f_A(f_B(\max(B_i))))))$. By computing $\max(B_i)$, f_B , and f_A , each agent can determine what social state gives him the highest level of utility. For individual 1 that state is σ_1 yielding a payoff value of 16. Individual 1 then must associate σ_1 with its corresponding strategy n -tuple by computing the inverse function $h^{-1}: Z \rightarrow SX$. The value of the function h^{-1} for σ_1 is s_1 . Individual 1 determines his action in strategy n -tuple s_1 by computing the function g_1^1 which yields the output a_1^1 .

Individual 1 will not conclude at this point that he should take action a_1^1 . He must first determine how many strategy n -tuples involve the action a_1^1 and determine which action every other agent will take to achieve their highest payoff. Since for a set of individuals $I = \{1, \dots, n\}$, there are m^n strategy k -tuples for each action $a_n^i \in S^i = \{a_1^i, \dots, a_m^i\}$, there are $m \times n-1$ strategy k -tuples in which a particular action of one individual is paired with I 's action. Since, in this example, there are two individuals and for each individual j , $S^j = \{a_1^j, a_2^j, a_3^j\}$, individual 1 must therefore consider three different strategy k -tuples.

Individual 1 reasons as follows: suppose I take action a_1^1 . Then individual 2, knowing my utility function and so on, will not take action a_1^2 yielding a payoff of 8, but will take a_3^2 yielding a payoff of 14, knowing as I do that he is instrumentally rational and so on.

But since I know that 2 knows what he knows, I should take action a_3^1 . But if I take action a_3^1 , then 2 will take action a_1^2 . If 2 will take action a_1^2 , then I will take action a_1^1 .

Individual 1 has completed the first round of an infinite loop. Regardless of which action either agent begins with, each encounters a potentially infinite loop.

Given these utility functions, there is no social state s_m such that the value of the function, $h^{-1}(f_A(f_B(\max(B_i))))$, for some individual i equals the value of the function, $h^{-1}(f_A(f_B(\max(B_j))))$, for every individual $j \neq i$. Therefore, no individual can compute his decision function. Hence, there exists a set of decision functions which are not effectively computable.

Perfect Information is not a sufficient condition for Pareto optimal equilibrium allocations.

Since there exists a case in which no decision function is effectively computable, perfect information is not sufficient to secure price-taking behavior. It follows that perfect information is not sufficient to achieve Pareto optimal outcomes. Agents are not themselves Turing Machines, they only possess Turing Machine computing capabilities. Since agents are capable of second-order awareness, if they discover a repeating loop in their computations, they will shut down until such time as new inputs change the structure of the situation.¹² Therefore, as we might expect, we can construct an allocation which consists of the set of initial endowments in social state s_0 and which yields a payoff vector, 6,11. Had each agent taken the "trade" option, they would have achieved an equilibrium allocation under social state s_5 in which one of them would have been better off than in social state s_0 and none worse off. Thus, social state s_5 is Pareto superior to social state s_0 . In this economy, in which moral normative constraints are absent but all of the conditions of standard

¹² Also, we can suppose each agent is "programmed" to recognize potentially infinite loops. This supposition is not subject to the results of the *Halting Problem for Turing Machines*. The *Halting Problem* shows that there is a function which is not Turing computable. A function h is defined so that, given any of a set of Turing Machine programs associated with a set of functions only some of which are computable, it would determine whether or not the program will halt. There is no Turing Machine which can compute the function h . The inputs of a decision function on the other hand are each effectively computable; thus, each are Turing-computable. When computing a set of inputs such as the values of the utility functions associated with *Table 1*, a Turing Machine will either halt in some non-standard position and fail to give an output, or it will enter a loop from which there is no exit. A Turing Machine can be programmed to identify potentially infinite loops in decision functions because decisions functions involve recursion from a double basis. Either way, we may suppose that the computational process ends.

accounts of a perfectly competitive market hold, the equilibrium allocation that individuals end up with is not Pareto optimal.

It may be objected that the result holds only in conjunction with the particular version of the Walrasian price mechanism presented here. In response, we claim that the result is not dependent on any particular price mechanism. The result holds even in the Arrow-Debreu model, where no action is taken until a message equilibrium is found. Given the utility functions described in the example, no message equilibrium is possible. Therefore, the result holds wherever agents must make simultaneous decisions. A model which involves simultaneous decisions is more realistic. Moreover, as long as agents must decide simultaneously, it is impossible to improve the situation with any type of information revealing mechanism.

Moral normative constraints are necessary for economically efficient outcomes of market interaction

Agents act as price-takers in versions of the First Welfare Theorem which specify moral normative constraints, because moral normative constraints ensure that agents act competitively. That is, moral normative constraints must be among the background assumptions of the First Welfare Theorem.

In order to facilitate the discussion it might help to see the logical form of our critique. Let “ $\{p_1 \& \dots, \& p_9\}$ ” symbolize the initial set of assumptions of the First Welfare Theorem in which agents are strict rational egoists and in which morality plays no role. Let “ $\forall x [EA_x \Rightarrow PO_x]$ ” symbolize the Theorem itself: *Every equilibrium allocation is Pareto optimal*. Our critique of the idea of a perfectly competitive market without moral normative constraints has the following logical form. We have supposed (S1):

$$(S1) \text{ Suppose: } \{p_1 \& \dots, \& p_9\} \Leftrightarrow \forall x [EA_x \Rightarrow PO_x]$$

But then we have shown (S2) which is a denial of (S1):

$$(S2) \text{ But: } \{p_1 \& \dots, \& p_9\} \Rightarrow \exists x [EA_x \& \neg PO_x]$$

We must conclude that the assumption set is not sufficient to achieve economically efficient outcomes of trade:

$$(S3) \text{ Thus, } \neg\{\{p_1 \& \dots, \& p_9\} \Rightarrow \forall x [EA_x \Rightarrow PO_x]\}$$

It follows that if every equilibrium allocation of commodities achieved through market interaction are Pareto optimal, at least one assumption in the initial set must be false. We must show which assumptions are false and specify alternative assumptions which are both necessary and sufficient for Pareto optimal equilibrium allocations of commodities achieved through trade.

Note that perfect information regarding every economically relevant variable is not itself a necessary condition. If we reduce the extent of perfect information, making it complete for some aspects of the action arena, but not for all, we have shown that agents will compute parametric decision functions, but reducing the extent of perfect information means that individuals' information will be asymmetrically-distributed for some economically-relevant variables. However, it is well known that asymmetric information makes force and fraud possible by creating the appropriate incentives. Therefore, in order for the First Welfare Theorem to hold in this economy when we eliminate (or restrict the range of) perfect information, we must also eliminate the incentives to take detrimental actions. Thus, perfect information regarding every economically relevant variable is neither sufficient nor necessary. I claim moral normative constraints are necessary and the assumption set which includes them is sufficient.

To make the case, I return once more to the question of why agents might act as price-takers. We now see that perfect information is not sufficient to secure price-taking behavior. If we can show in our model just how moral normative constraints can achieve what perfect information cannot, we will have demonstrated¹³ that moral normative constraints are necessary conditions of the First Welfare Theorem. A moral normative constraint can be construed as a moral right coupled with an incentive to comply with its demands, thus effectively restricting the types of actions agents can take.¹⁴ In principle, then,

¹³ In the first argument given in chapter two the *perfect information* component of assumption (p₆) did not extend to every individual's preference relation and natural strategy domain. In this argument, that is, in social situations defined by assumptions (p₁) - (p₉), agents have maximal information-processing capabilities, perfect information regarding every economically-relevant variable in a context that requires simultaneous, rather than sequential decisions in conjunction with the other conditions cited in assumption (p₆).

¹⁴ This is an important issue. There exists a debate between two broad types of conceptualizations of rights: a social choice conceptualization and a game-form conceptualization. Within each type there are variations, as one might expect. Regardless, for our purposes, we note the vital significance of the difference between a conceptualization which conceives of rights as *effectivity functions over social states* and one which depicts rights as *restrictions on an agents' strategy profiles*. We will adopt the second

it is possible to convert this example into a model in which the First Welfare Theorem holds by introducing moral rights (which render the decision function quasi-parametric in that some feasible actions by other agents which would have been required to be taken into consideration need not be now) and by altering the *Agents* subset to include a sufficient internal incentive to comply with all moral rights. Thus, in our example, we remove actions a_1^i and a_3^i for each individual i by introducing a moral right to true information and a sufficient internal incentive to comply. Each agent's decision function is then computable and each agent decides to trade.¹⁵

Since we have shown that Strict Rational Egoists will not comply with any enforcement mechanism, we must also change the *Agent* subset of the model and build in a sufficient *internal* incentive to comply with moral rights. We will discuss this point further in Chapter Six when we specify the moral conditions of economic efficiency. For now, all we need to see is that moral normative constraints can achieve what perfect information cannot. Therefore, moral normative constraints are missing conditions of the First Welfare Theorem.

Moral normative constraints as I have presented them involve both a moral rule and a sufficient internal incentive to comply with the moral rules. Both aspects are necessary. If we altered the agent set so that agents prefer moral visions of society or widely-accepted moral virtues such as honesty and if we did not also introduce moral rules, agents may not achieve efficient outcomes of market interaction. As I will explain in detail in chapter six, agents must hold in common beliefs about which types of behaviors are required and which are prohibited in order to achieve the required social ordering of behavior. It is not sufficient if each agent is free to work toward whatever moral vision of society or whatever moral virtues she "prefers". Agents must also hold in common beliefs about how to hold each other responsible.

type of conception. See van Hees (1995), Sen (1992), Gaertner *et al.* (1992), Sugden (1994).

¹⁵ It makes no difference that, in effect, individuals had no other option, because the very purpose of enforced rights is to restrict the range of options.

II. Discussion.

In 1776 Adam Smith shaped the world economy and consciousness to this day by claiming that when each person pursues his or her own interests they are together led as if by an *Invisible Hand* to achieve the common good. Introductory textbooks in economics represent this claim by the idea of a *Perfectly Competitive Market*. More advanced textbooks represent it by the *First Fundamental Theorem of Welfare Economics* (or: “First Welfare Theorem”). However, few, if any, proofs of the First Welfare Theorem explicitly specify the morality whose effects they presuppose. The First Welfare Theorem and its assumptions regarding agents has served as a point of departure for legal theory, appellate court decisions, economic analysis, and moral philosophy in the last two decades of the twentieth century. Tragically, many judges, public policy makers, economists and teachers of economics, and ordinary people take Smith’s *Invisible Hand* claim to mean that morality has no place in economic interaction as long as one follows the co-ordination conventions of a price mechanism. This paper summarizes a demonstration that a population of such selfish people cannot be guaranteed to achieve the common economic good in the absence of moral constraints on their behavior. More precisely, it shows that efficient outcomes of market interaction cannot be achieved without morality and then outlines a set of normative conditions which make efficient outcomes of trade possible—which includes an internal incentive to comply. This is where a Biblical view of property and stewardship play a central role.

Here is my argument: (1) *The First Welfare Theorem* presupposes the absence of deleterious effects of the actions of instrumentally rational, perfectly selfish agents. (2) Perfectly selfish agents cannot achieve efficiency in the absence of moral normative constraints, because (a) a presumption against non-market action entails a contradiction, (b) perfectly selfish agents will not behave competitively under the conditions of a *Perfectly Competitive Market* and (c) even given perfect information and perfect rationality efficient outcomes of trade are not guaranteed. Furthermore, (3) perfectly selfish agents will not collectively develop a system of self-enforcing rules and (4) perfectly selfish agents will not institute means for internalizing externalities. Therefore, in the absence of moral normative constraints, perfectly selfish agents pursuing their interests cannot achieve the “common good”. The common understanding of Adam Smith is wrong. Furthermore, Christians should repudiate the attitudes and practices of rational egoist capitalism recovering a clear sense of the Biblical view of property and the responsibilities and motives of a truly Christian stewardship.