Economies with Non-Convex Production and Complexity Equilibria

Christos H. Papadimitriou* Christopher A. Wilkens*

Abstract

The convexity assumptions required for the Arrow-Debreu theorem are reasonable and realistic for preferences; however, they are highly problematic for production because they rule out economies of scale. We take a complexity-theoretic look at economies with non-convex production. It had been known that in such markets equilibrium prices may not exist; we show that it is an intractable problem to achieve Pareto efficiency, the fundamental objective achieved equilibrium prices. The same is true for core efficiency or any one of an array of concepts of stability, with the degree of intractability ranging from $\mathrm{F}\Delta_2^P$-completeness to PSPACE-completeness. We also identify a novel phenomenon that we call *complexity equilibrium* in which agents quiesce, not because there is no way for any one of group of them to improve their situation, but because discovering the changes necessary for (individual or group) improvement is intractable. In fact, we exhibit a somewhat natural distribution of economies that gives an average-case hard complexity equilibrium.

*Computer Science Division, University of California at Berkeley, CA, 94720 {christos,cwilkens}@cs.berkeley.edu

1 Introduction

General Equilibrium Theory studies stable outcomes in markets — outcomes where each agent is doing as well as he can given the actions of others [17]. In the standard model, a market consists of consumers with initial endowments (vectors of goods) and preferences (utility functions), and firms with production sets (specifying what vector(s) of goods can be produced with each combination of raw materials); consumers own shares in firms.

By far the most studied kind of stable market outcome is the price equilibrium: Each firm optimizes its profit at market prices, and each consumer optimizes her utility at the same prices, selling her endowment and purchasing her preferred bundle of goods. Magically, this uncoordinated activity *clears the market:* no goods are left unsold, and all demand is satisfied. Most importantly, the resulting allocation of goods is *efficient* in the sense of Pareto: there is no allocation that is better in the sense that it dominates, in terms of utility, the allocation achieved via the price mechanism (this is known as the First Theorem of Welfare Economics).

Price equilibria had been studied by economists since the mid 19th century, but it was not until 1954 that Arrow and Debreu [3] made the idea irresistibly powerful and attractive by proving that (under assumptions) an equilibrium price vector is guaranteed to always exist. This result promises a kind of *Arrow-Debreu paradise*, where equilibrium is both beneficient (it achieves Pareto efficiency) and universally guaranteed[1]. The theory has spawned an entire area of Economics, and of course more recently a variety of results in Algorithmic Game Theory, including many algorithms for special cases (see [18], Chapters 5 and 6 for a survey, as well as [13, 11, 12] for production-specific algorithms).

There are of course wrinkles in General Equilibrium Theory. The existence proof in [3] is non-constructive, and this has been shown to imply some form of intractability, weaker than NP-completeness [19, 7]. The basic theorem holds for a very simplified model; in more realistic models parameters may be stochastic, time-varying, and generation-specific, among many other complications, and much work has been done addressing such difficulties. The model also hides tricky externalities (for example, production, or consumption, by one can harm the environment for all). Many other objections (e.g., that goods are available at different places and times) can be absorbed in the model by enlarging it. The focus of this paper is one of the most fundamentally objectionable assumptions of the theory, namely *the convexity assumption for production.*[2]

Convexity in utilities is quite natural: it states that you may draw less pleasure from your tenth evening dress than you did from your first. In contrast, convexity in production is very questionable because it *rules out economies of scale.* In other words, producing the hundredth airplane cannot, in any way, be easier than producing the first one. In the absence of this utterly unrealistic assumption — that is to say, in realistic economies — a price equilibrium may not exist, and thus the First Theorem cannot guarantee Pareto efficiency: Paradise lost.

Market Equilibrium Theory without Convexity of Production. Since price equilibria may not exist in the absence of convexity in production, economists have studied the set of Pareto optima (which do generally exist). The first work in this line was by Guesnerie [9], whose stated goal was "to characterize precisely Pareto-optimal states and to examine the possibility of achieving them in a decentralized economy" (a task which is, as we point out, unattainable for reasons of

[1] In an excellent Microeconomics textbook [16] one reads after the statement of the First Theorem: "You should now be hearing choirs of angels and choruses of trumpets. The invisible hand of the price mechanism produces equilibria that cannot be improved upon." The author goes on to expose and discuss the many problems of the theory.

[2] Convexity in production refers to convexity of the set of net production vectors. For example, if a firm can produce according to net production (input/output) vectors y_1 and y_2, then it can also produce according to any production vector $\alpha y_1 + (1-\alpha)y_2$ for $\alpha \in [0,1]$.

complexity). Since then, a vast literature has developed; for two excellent surveys, see [6, 1] — the first one actually contains a discussion of computation.

A standard approach is to assume that firms price goods at marginal cost, in other words, postulate that prices depend on production decisions; this assumption is quite strong and rather artificial and unrealistic, but it often yields an allocation that is Pareto efficient — not always, of course. The state of the art in this direction (e.g. the marginal pricing rule of [5]) still seeks a decentralized model of agent behavior that is guaranteed to achieve Pareto efficiency. Our results, outlined next, suggest that there are huge computational impediments in the way of this ambition.

Our Results: Computational Complexity in Markets with Non-Convex Production. We study markets with non-convex production from the perspective of computational complexity. To the best of our knowledge, the only other work with a computational flavor is [21], which employs dimensional communication complexity to differentiate this case from the convex one: it is shown in [21] that $\approx m \cdot n$ reals are needed to achieve Pareto efficiency in this case, where m is the number of goods and n the number of agents and firms, as opposed to only m in the convex case.

We show that the theory of markets with non-convex productions is plagued with very bleak negative complexity results, as many natural concepts of rationality are hard for various levels of the polynomial hierarchy. We start by showing that computing a Pareto efficient outcome in a market with non-convex production is $\mathrm{F}\Delta_2^P$-hard. Economists regard Pareto efficiency as a sine qua non for any concept of stability or rationality in markets. Hence, our negative result for the complexity of finding Pareto efficient outcomes is a lower-bounds for any "reasonable" equilibrium concept. Finally, in sections 4 and 5, we give similar results for two concepts of stability more sophisticated than Pareto efficiency: It is $\mathrm{F}\Sigma_2^P$-hard to tell if an allocation is in the core (no coalition of agents has an incentive to defect and create its own economy). And for a natural models of sequential production, we show that computing equilibria is $\mathrm{F}\Delta_3^P$-hard and PSPACE-hard, respectively.

Perhaps most significantly, we show in the process that such economies can have a novel kind of "equilibrium," from which deviation may yield tremendous improvement for any and all agents, but the agents are stuck at a suboptimal solution of a particular instance of an NP-hard optimization problem. We call such a situation a *complexity equilibrium* (Definition 6.2). When agents are at such an equilibrium, standard complexity-theoretic assumptions imply that no computationally efficient procedure would generally allow them to improve – indeed, it is even intractable to recognize that improvement is possible. With the exception of deliberate complexity-theoretic studies in game theory (e.g. [20]), we are not aware of other natural economic situations in which computational complexity begets stability. Interestingly, we also present a somewhat natural *average-case* NP-hard construction of a complexity equilibrium.

2 Foundations and Models

In this section we introduce the standard economic model and relevant complexity classes.

2.1 The Economic Model

We will employ a slightly simplified version of the standard private ownership economy used in general equilibrium theory [17]. We define an economy E as follows:

Agents: An economy has n agents.

Goods: An economy has m divisible, tradable goods.

Utilities: Each agent i has a utility function $u_i : \mathbb{R}^m \to \mathbb{R}$ mapping bundles of goods to amounts of utility. An agent's consumption in an economy is specified by a vector of goods $x_i \in \mathbb{R}^m$.

Our reductions use a simple form of utilities known as *Leontief* utilities, which take the form

$$u(x) = \min_j \frac{x_j}{\alpha_j} \ ,$$

i.e. goods are demanded in constant proportions specified by the parameters $\{\alpha_j\}$ (possibly 0).

It is reasonable to make the following assumptions about utility functions (see [22] for a discussion of their absence):

1. Given x, it is possible to obtain $u_i(x)$ in polynomial time (e.g. $u_i(x)$ is efficiently computable or we are given an oracle for $u_i(x)$).
2. Each function $u_i(x)$ satisfies standard convexity assumptions as stated by Arrow and Debreu [3] (continuity, convexity of the upper level set, and nonsatiation).

Endowments: Each agent i is endowed with a quantity of each good, i.e. a vector $e_i \in \mathbb{R}^m$.

Production: Each agent in an economy owns[3] a set of production units $F_i = \{f_k : \mathbb{R}^m \to \mathbb{R}^m\}$ where f_k maps bundles of input goods to bundles of output goods. The behavior of a production firm is often specified by a net production vector $y_k = f_k(x_k) - x_k$, where x_k is the vector of goods consumed by firm k (note that this is different from x_i, the vector of goods consumed by agent i).[4]

Our reductions use one very simple form of non-convex production function, namely *Leontief production functions with fixed costs.* Such a function f takes the form

$$f(x) = z \cdot \max\left(\min_j \left(\frac{x_j - \beta_j}{\alpha_j}\right), 0\right)$$

where z is a bundle of goods and β_j is a fixed cost of each good required to have positive production. Interestingly, the addition of fixed costs is sufficient to force the agent to solve a discretized problem. This will be a key technique in our reductions.

We make the following assumptions about production in the economy:

1. A production function $f(x)$ is efficiently computable for all x.
2. The total set of production possibilities is closed and bounded. The total set of production possibilities contains any net production vector that may achieved by the economy. I.e. it contains a vector $y \geq 0$ if and only if there is a set of vectors $\{y_k \in Y_k\}$ such that $\sum_i e_i + \sum_k y_k = y$.
3. There is no x such that $f(x) - x \geq 0$ other than $x = 0$ (no free lunch).
4. For all $f(x)$, a bounded input x implies a bounded output $f(x)$.

With the exception of the computability assumption, these are standard or weakened versions of standard assumptions in the economic literature. The efficient computability assumption is nonstandard insofar as the issue has not been considered.

Finally, we recognize that smoothness is a common assumption in economics. While the functional forms we use for f (and u) are not smooth, they may be made smooth without affecting the results in this paper.

[3] In the standard private ownership economy, agents are said to own *shares* in production firms and receive the appropriate fraction of the profit. Since we avoid discussion of prices, the total ownership restriction avoids the issue of cooperative production for which there is already a literature, e.g. [8].

[4] Standard General Equilibrium Theory specifies a production firm by a set of possible net production vectors Y_k instead of a function f_k. While there are scenarios that differentiate the two representations, we will not encounter them here.

We also use the following standard economic vocabulary (see [17]):

Definition 2.1 *An* economic allocation *(hereafter an allocation) is an assignment* (x, y) *such that* x_i *is the vector of goods consumed by agent* i *and* y_k *is the vector of net outputs for firm* k*.*

An allocation is feasible *if the amount consumed is less than the amount available in the economy, i.e.*

$$\sum_i x_i \leq \sum_i e_i + \sum_k y_k$$

where $y_k = f_k(x_k) - x_k$ *for the vector of goods* x_k *used as input to* f_k*.*

Definition 2.2 *An economic allocation* (x_1, y_1) *is said to be* strictly Pareto preferred *to another allocation* (x_2, y_2) *if some agent receives more utility in allocation 1 than in allocation 2 and no agent receives less utility.*

An allocation (x^*, y^*) *is* Pareto efficient *or* Pareto optimal *if no feasible allocation is strictly Pareto preferred to it.*

Definition 2.3 *The* social welfare W *of an allocation* (x, y) *is the sum of the utilities obtained by consumers in the economy, i.e.*

$$W = \sum_i u_i(x_i).$$

The particular form of Leontief utilities gives the following:

Proposition 2.4 *(Free disposal under Leontief Utilities.) When utility functions are Leontief, then* $x' \geq x$ *implies* $u_i(x') \geq u_i(x)$*.*

This allows us to make a key assumption: when an agent has a good $\hat{s}$ and only one possible use for that good, we may assume that $\hat{s}$ is applied to that use because no harm can be done.

2.2 The Polynomial Hierarchy

Our computational complexity results will locate variations on the General Equilibrium problem in different classes of the polynomial hierarchy. The relevant portions of the polynomial hierarchy, Σ_k^P and Δ_k^P, are defined recursively as

$$\Sigma_0^P = \Delta_0^P = \mathrm{P}$$

$$\Sigma_k^P = \mathrm{NP}^{\Sigma_{k-1}^P}$$

$$\Delta_k^P = \mathrm{P}^{\Sigma_{k-1}^P}$$

in other words, Σ_k^P is equal to NP with an oracle for Σ_{k-1}^P. A prefixed "F" denotes the corresponding class of functional problems, e.g. $\mathrm{F}\Sigma_k^P$.

Krentel [14] defines the related class OptP as the class of problems that may be expressed as the maximum (or minimum) value along any branch of a nondeterministic Turing machine. Relevant to our work, he shows that any OptP-complete problem is complete for the class $\mathrm{F}\Delta_2^P$ and shows a similar generalization to $\mathrm{F}\Delta_k^P$ [15]. For our purposes, we use the fact that any OptP problem is in $\mathrm{F}\Delta_2^P$.

3 Computing Pareto Optima

In Economics, Pareto efficiency (the requirement that there be no allocation preferred to the current one by all) is essentially a prerequisite for any reasonable solution concept or prediction. By showing negative complexity results for finding a Pareto efficient allocation, we lower bound the complexity of any equilibrium concept that achieves Pareto efficiency. Our first result classifies the hardness of computing a Pareto efficient allocation:

Theorem 3.1 *Computing a Pareto efficient allocation in an economy with non-convex production functions and polynomial-time computable utility functions is* $\mathrm{F}\Delta_2^P$*-complete.*

We prove this theorem after introducing two gadgets.

3.1 Gadgets

The proofs for Theorems 3.1 and 6.3 will construct economies out of the following gadgets:

The Choice Gadget. The choice gadget $Choice(\alpha_1 \cdot \hat{j}_1, \ldots \alpha_c \cdot \hat{j}_c)$ enforces indivisibility: given a set of production options, agent i must choose exactly one of the possible output goods. Specifically, when agent i has a choice gadget $Choice(\cdots)$, the economy includes the following:

Goods: $\hat{s}$, $\hat{j}_1, \ldots \hat{j}_c$.

Production: agent i owns firms with the following production functions:

$$\forall \hat{j}_k : \quad x_{\hat{j}_k} = f_{j_k}(x_{\hat{s}}) = \alpha_k \cdot \max(x_{\hat{s}} - 1, 0) \ .$$

Endowment: $e_{i,\hat{s}} = 2$.

For each good $\hat{j}_k$, agent i has a production function to turn $x_{\hat{s}}$ units of good $\hat{s}$ into $\alpha_k \cdot (x_{\hat{s}} - 1)$ units of good $\hat{j}_k$. Since agent i has only 2 units of good $\hat{s}$, it follows that in any allocation, only one good $\hat{j}_k$ may be produced in positive quantities.

In order to ensure that all of good $\hat{s}$ is consumed, we stipulate that no agent has any other use for good $\hat{s}$, either as a source of utility in consumption or as an input to production. Thus, by Proposition 2.4 (free disposal), we may assume that agent i will use all 2 units of good $\hat{s}$ in equilibrium and, therefore, produce exactly α_k units of the chosen good $\hat{j}_k$.

The Limit Gadget. The limit gadget enables the economy to limit the production of a specific good $\hat{j}$ to α units. An instance of $Limit(\hat{j}, \alpha)$ consists of

Goods: $\hat{\hat{j}}$, $\hat{r}$, and $\hat{j}$.

Production: (owned by agent i):

$$x_{\hat{j}} = f(x_{\hat{\hat{j}}}, x_{\hat{r}}) = \min(x_{\hat{\hat{j}}}, x_{\hat{r}}) \ .$$

Endowment: $e_{i,\hat{r}} = \alpha$.

The good $\hat{r}$ acts as a limiting reagent in the production function $f(x_{\hat{\hat{j}}}, x_{\hat{r}})$ — since the endowment of $\hat{r}$ is fixed at α, agent i may produce as much $\hat{j}$ as desired up to α units.

To enforce this limit, all production functions that produce $\hat{j}$ are modified to produce $\hat{\hat{j}}$, thereby forcing all of good $\hat{j}$ to come from $f(x_{\hat{\hat{j}}}, x_{\hat{r}})$ or an endowment.

3.2 Proof

First, we must note that Pareto optima always exist in our economies:

Proposition 3.2 *Under the assumptions of Section 2.1, a Pareto efficient allocation always exists.*

Proof: Observe that if allocation (x', y') is strictly Pareto preferred to (x, y), then (x', y') must have a higher social welfare than (x, y). Thus, since the maximum social welfare W^* is well defined, any allocation that achieves W^* is an allocation for which no strictly Pareto preferred allocation exists, i.e. a Pareto optimum. □

We now prove Theorem 3.1.

Proof: (Proof of Theorem 3.1.) First, we show that an efficient allocation may be computed in $\mathrm{F}\Delta_2^P$.

Consider the following problem: *compute a feasible allocation* (x, y) *with social welfare at least* W. It may be solved in NP by guessing the goods x_i consumed by each agent and the goods x_k used as inputs by each firm. (The assumptions of computability imply that we may efficiently compute (x, y) and W from the x_i and x_k vectors.) As in Proposition 3.2, an allocation with optimal social welfare W^* must be Pareto efficient. Thus, a Pareto efficient allocation may be expressed as the optimum of an NP problem, so it is in OptP and therefore $\mathrm{F}\Delta_2^P$.

To show that computing a Pareto efficient allocation is $\mathrm{F}\Delta_2^P$-hard, we reduce the $\mathrm{F}\Delta_2^P$-complete problem Weighted MAX-SAT [14] to an economy. Let $(\Phi = \bigwedge_j \phi_j, \{\alpha_j\})$ be a Weighted MAX-SAT instance, i.e. we desire a boolean assignment χ to the CNF formula Φ that maximizes $\sum_j \alpha_j \phi_j(\chi)$. Consider the following economy:

Agents: One agent i for each variable χ_i and one agent j for each clause ϕ_j.

Goods: A utility good $\hat{\gamma}$.

For each SAT variable χ_i: two goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$.

For each clause ϕ_j: one good $\hat{\phi}_j$.

Utilities: $u_i(x) = x_{\hat{\gamma}}$.

Production: Each variable agent i owns:

$$Choice(\hat{\chi}_i, \hat{\bar{\chi}}_i) \ ,$$

and each clause agent j owns:

$$\forall \chi_i \in \phi_j : \ \ x_{\hat{\phi}_j} = f_{\phi_j,\chi_i}(x_{\hat{\chi}_i}) = x_{\hat{\chi}_i}$$

$$\forall \bar{\chi}_i \in \phi_j : \ \ x_{\hat{\phi}_j} = f_{\phi_j,\bar{\chi}_i}(x_{\hat{\bar{\chi}}_i}) = x_{\hat{\bar{\chi}}_i}$$

$$Limit(\hat{\phi}_j, 1)$$

$$x_{\hat{\gamma}} = f_{\gamma,\phi_j}(x_{\hat{\phi}_j}) = \alpha_j \cdot x_{\hat{\phi}_j} \ .$$

In this economy, a positive quantity of good $\hat{\chi}_i$ (respectively $\hat{\bar{\chi}}_i$) signifies that χ_i was set to true (respectively false). The choice gadget ensures that only one good ($\hat{\chi}_i$ or $\hat{\bar{\chi}}_i$) may be present in positive quantities. Similarly, a positive quantity of good $\hat{\phi}_j$ signifies that clause ϕ_j was satisfied, enforced by the production functions f_{ϕ_j,χ_i} and $f_{\phi_j,\bar{\chi}_i}$. Finally, the clause agents turn $\hat{\phi}_j$ into a utility good $\hat{\gamma}$ using f_{γ,ϕ_j}. The weighted sum of the satisfied clauses will be the amount of $\hat{\gamma}$ in the economy. The limit gadget ensures that each clause is only "counted" once.

Since agents only desire $\hat{\gamma}$, it follows that any Pareto efficient allocation must create the maximum amount of $\hat{\gamma}$, i.e. solve the weighted MAX-SAT instance. Thus, computing a Pareto efficient allocation is $\mathrm{F}\Delta_2^P$-complete. □

Of course some non-convex economies may have price equilibria, or equilibria in marginal prices; is the previous complexity result irrelevant in such favorable special cases? Alas, we show below that these cases are hard to recognize, and even if we know our economy is such, computing these prices is intractable:

Theorem 3.3 *It is NP-complete to distinguish between economies that have no price equilibria or equilibria in marginal prices, and those that have both.*

A sketch is in the Appendix.

4 Computing Core Allocations

Pareto efficient allocations are stable in a cooperative sense, that is, with respect to a concept of deviation that requires all agents to cooperate and change their production and consumption. in this section, we consider allocations that are also stable with respect to certain selfish defections. A standard concept of rationality in economics is the *core*. An allocation is said to be in the core if no coalition would prefer to defect, i.e. no subset of agents can achieve strictly higher utility among themselves by creating a separate economy in which they are the only agents [17]. Price equilibria are always in the core; however, since price equilibria do not exist as such, we potentially lose this property when we lose convexity. A complexity equilibrium therefore is an allocation from which it is intractable to find another allocation where some agents do strictly better using only their own endowment and production technologies.

We find that computing core allocations is harder than computing Pareto optima:

Theorem 4.1 *Computing an allocation in the core is* $\mathrm{F}\Sigma_2^P$*-complete (and such an allocation may not exist).*

Additionally, if we relax our rationality requirements to include only single-agent deviations, then rational allocations are guaranteed to exist and are as easy to find as Pareto optima:

Theorem 4.2 *Computing an allocations that is rational with respect to single-agent deviations and Pareto improvement is* $\mathrm{F}\Delta_2^P$*-complete.*

The remainder of Section 4 contains proof sketches for Theorems 4.1 and 4.2. Certain full proofs may be found in Appendix A.

4.1 Proof Sketches

We use new gadgets; however, we only sketch the relevant one here (the complete definitions can be found in Appendix A).

The Circle-of-Death Gadget. The circle-of-death $COD(\hat{d})$ constructs a group of three agents for which it is impossible to achieve a core allocation unless they may consume $\geq \frac{1}{2}$ units of good $\hat{d}$. (The construction is given in Appendix A.)

Proof: *(Proof sketch for Theorem 4.1.)* For $\mathrm{F}\Sigma_2^P$-completeness, the difficult direction is to show that it is $\mathrm{F}\Sigma_2^P$-hard to find a core allocation (x^c, y^c). We reduce from the $\mathrm{F}\Sigma_2^P$-complete problem Σ_2SAT [2]: Σ_2SAT [2]:

$$\exists x_1 \forall x_2 \Phi(x_1, x_2) \ .$$

(The full reduction is given in Appendix B.1.) It is easy to construct an economy that attempts to satisfy Φ. The trick is to engineer the payoffs such that a coalition defects if there is an x_2 that falsifies Φ given x_1. We encode x_1 in the identities of the agents who defect so the defecting coalition can only produce one value for x_1, and we use a circle of death gadget to ensure that no allocation intentionally falsifying Φ can ever be in the core. As a result, any allocation in the core must make Φ true for any setting of x_2, i.e. x_1 is as desired.

□

Proof: *(Proof sketch for Theorem 4.2.)* To show $\mathrm{F}\Delta_2^P$-hardness, first, assume a single agent i to leaves the economy. The defecting agent will maximize his utility using only his endowment e_i and his own production functions F_i. We will call his resulting utility his batna,[5] b_i. Note that we can compute b_i in $\mathrm{F}\Delta_2^P$, and, given the b_i's, optimizing production subject to individual rationality is also in $\mathrm{F}\Delta_2^P$. □

5 Computing Equilibria in a Sequential Model

Thus far, we have considered the most common models of rationality in markets — individual defection and core rationality. However, since production often involves intermediate goods, it seems natural to consider a model that makes this sequential property explicit.

For example, consider a production sequence in which $\hat{a}$ is transformed into $\hat{b}$, $\hat{b}$ is transformed into $\hat{c}$, and $\hat{c}$ is consumed. Moreover, imagine that the agent who transforms $\hat{b}$ into $\hat{c}$ would rather just keep the $\hat{b}$ that he gets. It seems natural that this agent should be able to defect after receiving $\hat{b}$, therefore disrupting future production. Core rationality ignores this possibility because it requires that defecting agents be excluded from the economy.

Since we are not aware of a standard, general model of sequential production, we adopt what we believe is a natural model. Specifically, we augment the production model to specify time:

Definition 5.1 *A* sequential production plan *is a specification* $(\{x_i\}, \{x_{k,t}\})$ *including the vector of goods* x_i *consumed by agent* i *and the vector of production inputs* $x_{k,t}$ *used by firm* k *at time* t*. A feasible production plan is one in which all production inputs required at time* t *exist in the economy, i.e.*

$$\sum_k x_{k,t} \leq \sum_{\tau=0}^{t-1} \sum_k \left(f_k(x_{k,\tau}) - x_{k,t}\right) + \sum_i e_i \ .$$

To ensure that all sequential production plans have polynomial size, we require that each function f_k may be used at most once and that some production function f_k is used at each timestep.

First, we assume that defectors are isolated for the remainder of the production plan, i.e. an agent will participate until time t_d and then choose to deviate, after which he cannot trade and can only use his own production technologies F_i. In this setting, we classify the complexity and show that complexity equilibria exist under the natural generalization:

Theorem 5.2 *In a model in which defecting individuals will face subsequent isolation, it is* $\mathrm{F}\Delta_3^P$*-complete to find the optimal sequential production plan in which no individual wishes to defect.*

[5]The acronym BATNA, meaning "Best Alternative To a Negotiated Agreement," is commonly used in negotiation theory.

Second, we observe that isolation is not always a credible threat. We would like a "subgame perfect" allocation; however, a subgame perfect equilibrium may require an exponentially large specification. Thus, we ask for a production plan that is consistent with the realization of some subgame perfect equilibrium. We show that this problem is computationally harder:

Theorem 5.3 *In a model in which defectors are not isolated, it is PSPACE-hard to find a sequential production plan that is consistent with a subgame perfect equilibrium.*

The proof for Theorem 5.2 may be found in Appendex B.2. The proof for Theorem 5.3 is omitted.

6 Complexity Equilibria

Perhaps the most interesting complexity-theoretic phenomenon in non-convex economies is the existence of allocations that are stable because profitable deviation is computationally intractable. To formalize complexity equilibria, one must recognize that stable outcomes in Economics are defined in terms of an appropriate concept of *deviation.*

Definition 6.1 *A deviation scheme is a mapping $\mathcal{D}$ assigning each feasible allocation and subset of agents a set of feasible allocations. Intuitively, if (x, y) is a feasible allocation and $S \subseteq [n]$, then $\mathcal{D}((x, y), S)$ is the set of all allocations to which the agents in S can drive the economy in one step called* a deviation by S. *Deviation $(x', y') \in \mathcal{D}((x, y), S)$ is a* profitable deviation by S *if each $i \in S$ has at least as good utility in (x', y') than in (x, y), and at least one agent $i \in S$ has strictly better utility. A $\mathcal{D}$-equilibrium is an allocation (x, y) such that for all S $\mathcal{D}((x, y), S)$ contains no profitable deviations.*

Suppose that, for all allocations (x, y) $\mathcal{D}((x, y), S)$ is the set of all feasible allocations whenever $S = [n]$, and is the empty set otherwise; then $\mathcal{D}$-equilibria are precisely the Pareto optimal allocations. To define the core, $\mathcal{D}((x, y), S)$ contains all feasible allocations that are also feasible if the endowments, consumption, and production by agents not in S is set to zero. And for the sequential production model, $\mathcal{D}((x, y), \{i\})$ contains all allocations that can be achieved by having agent i unilaterally change her production decisions; all other values of $\mathcal{D}$ are empty.

We can now define complexity equilibria:

Definition 6.2 *We say that a family of economies $\mathcal{E}$ has* complexity equilibria *with respect to deviation scheme $\mathcal{D}$ if the following propblem is NP-complete: Given an allocation $(x^+, y^+)_E$ in an economy $E \in \mathcal{E}$, find a profitable $\mathcal{D}$-deviation.*

Theorem 6.3 *Economies with non-convex production possibilities have complexity equilibria with respect to Pareto improvement, core rationality, and sequential rationality. Moreover, the inefficiency of the complexity equilibrium is unbounded.*

The proof of this theorem follows later in this section. Through standard complexity theory arguments involving so-called *complexity cores*, this result implies that, unless P = NP, there are infinite families of allocations on which any group of polynomial-time agents would be almost always stuck at a very inefficient allocation. Moreover, we show the existence of families of economies with complexity equilibria relative to average-case NP-hardness. We refer the reader to Bogdanov and Trevisan's survey [4] for background on average-case complexity.

Theorem 6.4 *There exists distribution of economies $(\mathcal{D}, \mathcal{E})$ with complexity equilibria from which improvement is average-case NP-hard.*

Significantly, the family of economies exhibiting this behavior is a somewhat natural model — certainly quite natural compared to many of the other distributional problems for which one can prove average-case results. We sketch the proof in Appendix C.

Proof: (Of Theorem 6.3.) Technically, this theorem follows from the construction in Theorem 3.1 and the fact that given an assignment satisfying k clauses of a CNF, it is NP-hard to satisfy $k+1$ clauses. However, this argument is quite nonconstructive. To be more explicit, we construct a simple family of economies $\mathcal{E}$ in which agents only receive a significant amount of utility if they find a solution to a CNF SAT formula Φ.

Let $\Phi = \bigwedge \phi_j$ be a CNF SAT instance with n_Φ variables χ_i and m_Φ clauses ϕ_j. We construct the following economy E_Φ:

Agents: One agent i for each SAT variable χ_i.

Goods: For each variable χ_i: two goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$.

For each clause ϕ_j: one good $\hat{\phi}_j$.

Utilities: $u_i(x) = m_\Phi \cdot \min_j x_{\hat{\phi}_j}$.

Production: Each agent i owns:

$$Choice(\hat{\chi}_i, \hat{\bar{\chi}}_i)$$

(this implicitly adds one source good $\hat{s}_i$ per variable χ_i),

$$\forall \phi_j \text{ s.t. } \chi_i \in \phi_j : \quad x_{\hat{\phi}_j} = f_{\phi_j, \chi_i}(x_{\hat{\chi}_i}) = x_{\hat{\chi}_i} \quad ,$$

$$\forall \phi_j \text{ s.t. } \bar{\chi}_i \in \phi_j : \quad x_{\hat{\phi}_j} = f_{\phi_j, \bar{\chi}_i}(x_{\hat{\bar{\chi}}_i}) = x_{\hat{\bar{\chi}}_i} \quad .$$

Each agent i also owns

$$\forall \phi_j : \quad x_{\hat{\phi}_j} = f\phi_j, \hat{s}_i(x_{\hat{s}_i}) = \frac{1}{(nm)^2} x_{\hat{s}_i}$$

This economy behaves similarly to the one in Theorem 3.1, except that it is based on a standard CNF SAT instance instead of a MAX-SAT instance. Again, agents are intended to choose $\hat{\chi}_i$ or $\hat{\bar{\chi}}_i$ to pick a setting of variable χ_i. However, since we want all clauses satisfied, we use a Leontief utility function to give utility if and only if all clause goods are present.

It immediately follows that agents may obtain positive utility from an allocation if and only if Φ is satisfied. When the agents do satisfy Φ, each variable may occur in at most m_Φ clauses and, therefore, it is always possible to produce at least $\frac{1}{m_\Phi}$ units of each clause good $\hat{\phi}_j$ and, consequently, attain a total social welfare W of at least 1.

Thus, the allocation in which the agents use all their $\hat{s}_i$ in $f\phi_j, \hat{s}_i(x_{\hat{s}_i})$ is a complexity equilibrium — it is certainly not Pareto efficient if Φ has a solution; however, finding an allocation that is Pareto preferred requires satisfying Φ, which is NP-hard. Moreover, if we eliminate $f\phi_j, \hat{s}_i(x_{\hat{s}_i})$ altogether, the relative inefficiency (relative social welfare of the complexity equilibrium compared to a true Pareto optimum) is unbounded because the social welfare at the complexity equilibrium is 0 and the social welfare at a Pareto optimum is at least 1.

□

7 Discussion and Open Problems

We showed that economies with nonconvexities — in other words, real economies — can be theaters of extreme complexity phenomena, including a novel kind of equilibrium in which agents quiesce because of the intractability of the task of finding a better allocation. One remark here is in order: economists often respond to complexity results such as the PPAD-completeness of Nash equilibria by questioning the relevance, and plausibility in real life, of the complex games with specialized structure that arise in those reductions. In the present situation, however, the intractability is, intuitively, more "generic." Nonconvex optimization is a hard problem, and in hard optimization problems "gaps" between optima and defaults are common. As a result, the present complexity results may be a little more compelling to economists.

One could hope for a proof that, in a well-defined sense to be determined, nonconvex economies are "often," or even "almost always," computationally hard. Our average-case hard construction takes a step in this direction, and, we believe, gives hope that stronger results are possible.

A Appendix: Gadgets

The SAT Gadget. The $SAT_\Phi(\{\hat{\chi}_i\}, \{\hat{\bar{\chi}}_i\}, \hat{\Phi}_{TRUE}, \hat{\Phi}_{FALSE})$ gadget enables an agent i to evaluate an arbitrary boolean formula Φ to produce exactly one unit of either $\hat{\Phi}_{TRUE}$ or $\hat{\Phi}_{FALSE}$. Let Φ be expressed as a tree T on which each literal in Φ is a leaf and each internal node represents the AND or OR of its children. Then the SAT gadget is described by the following subset of an economy:

Goods: For all variables χ_i in Φ: goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$.

For each internal node t in tree T: goods $\hat{t}$ and $\hat{\bar{t}}$.

Let r refer to the root of the tree. Then $\hat{r}$ and $\hat{\bar{r}}$ are synonyms for $\hat{\Phi}_{TRUE}$ and $\hat{\Phi}_{FALSE}$.

Production: For all AND nodes t in T with children c_j: the functions

$$x_{\hat{t}} = f_t(x) = \min_j x_{\hat{c}_j}$$

$$\forall \bar{c}_j : \quad x_{\hat{\bar{t}}} = f_{\bar{t}}(x) = x_{\hat{\bar{c}}_j}$$

and for all OR nodes t with children $c_1, \ldots$: two functions

$$\forall c_j : \quad x_{\hat{t}} = f_t(x) = x_{\hat{c}_j}$$

$$x_{\hat{\bar{t}}} = f_{\bar{t}}(x) = \min_j x_{\hat{\bar{c}}_j}$$

Finally, we ensure that at most one unit of the true and false good exists:

$$Limit(\hat{\Phi}_{TRUE}, 1)$$

$$Limit(\hat{\Phi}_{FALSE}, 1)$$

(When necessary, constant scalars may be added to ensure that at least one unit of $\hat{\Phi}_{TRUE}$ or $\hat{\Phi}_{FALSE}$ is created.)

In the manner of previous SAT reductions, this gadget allows direct evaluation of Φ given sufficient quantities of the goods setting each variable χ_i. The main differences from our previous reductions are that this gadget evaluates arbitrary formulas and that it explicitly signals false as well as true.

The Circle-of-Death Gadget. The circle-of-death constructs a group of three agents who will continually defect unless they are given a particular good. The gadget $COD(\hat{d})$ includes the following:

Agents: Three agents 1, 2, and 3.

Goods: For each agent i, there is one source good $\hat{s}_i$ and one product good $\hat{p}_i$. There is also a deactivator good $\hat{d}$.

Endowments: Agent i is endowed with 2 units of $\hat{s}_i$ and nothing else.

Utilities: Agent i has utility function $u_i(x) = x_{\hat{p}_i}$.

Production: Agent i has a production function for producing the *bundle*

$$[x_{\hat{p}_{i-1}}, x_{\hat{p}_i}] = f_i(x) = [2, 1] \cdot \max(\min(x_{\hat{s}_{i-1}} - 1, x_{\hat{s}_i} - 1), 0)$$

and a deactivated production function

$$x_{\hat{p}_i} = f_{\hat{d},i}(x_{\hat{s}_i}, x_{\hat{d}}) = \min(x_{\hat{s}_i}, 6x_{\hat{d}})$$

The f_i are designed with three properties:

1. Because of fixed costs, only one function f_i may be used at a time. Thus, if f_i is used, agent $i+1$ will get 0 units of utility, agent i will get 1 unit, and agent $i-1$ will get 2.
2. For any choice of f_i to use, the function f_{i+1} gives 2 units of utility to agent i, 1 unit to agent $i+1$ and 0 units to agent $i-1$. Moreover, agents i and $i+1$, both of whom strictly prefer using f_{i+1}, have both the endowment and production technology to achieve this result in isolation. Thus, in the absence of $\hat{d}$, there is always a defecting coalition.
3. Agents may opt out of the circle-of-death and use the deactivator good to produce utility if present, but the "losing" agent in the cycle must be able to generate at least 1 unit of utility if the cycle is to be broken, requiring $\geq \frac{1}{2}$ units of $\hat{d}$.

Consequently, if $< \frac{1}{2}$ units of $\hat{d}$ are available in the economy, then the core is empty, i.e. some pair of agents would always benefit from defecting.

Binary Counting Gadget. The gadget $BCG(\hat{\gamma}, \hat{j}_0, \ldots \hat{j}_c)$ treats $x_{\hat{j}_0}, \ldots x_{\hat{j}_c}$ as the binary representation of a c-bit number x and produces x units of $\hat{\gamma}$. It includes the following:

Goods: The "counting" good $\hat{\gamma}$ and the c input bit goods $\hat{j}_k$.

Production: For each good $\hat{j}_k$:

$$x_{\hat{\gamma}} = f_k(x_{\hat{j}_k}) = 2^k \cdot x_{\hat{j}_k} \quad .$$

Generalized Choice Gadget. The generalized choice gadget $GChoice(\hat{s}, x_1, \ldots x_c)$ is identical to the choice gadget (see Section 3) except that the input good $\hat{s}$ is provided by the economy and the choice is over bundles x_k instead of individual goods $\hat{j}_k$:

Goods: A source good $\hat{s}$ and bundles $x_1, \ldots x_k$ over the space of all other goods $\mathbb{R}^{m-1}$.

Production: Agent i owns firms with the following production functions:

$$\forall j_k : \quad z = f_{j_k}(x_{\hat{s}}) = x_k \cdot \max(x_{\hat{s}} - 1, 0) \quad .$$

$$Limit(\hat{s}, 2)$$

B Appendix: Rational Reductions

B.1 Core Rationality

Proof: *(of Theorem 4.1.)* It is easy to certify that an allocation (x, y) is not in the core (i.e. testing core membership is in coNP): demonstrate a coalition and an allocation (x', y') such that the coalition strictly prefers (x', y') to (x, y). Thus, a core allocation (x^c, y^c) may be found in $F\Sigma_2^P = FNP^{NP}$ as follows: guess the core allocation and use an oracle call to check that it is in the core.

Next, we give the reduction from the $F\Sigma_2^P$-complete problem Σ_2SAT [2]: find x_1 such that $\forall x_2 \Phi(x_1, x_2) = 1$. Let n_Φ be the number of variables in Φ.

We will describe the economy in stages. The globally relevant parts of the economy include:

Agents: Two directors A and B.

For each variable χ in x_1: two agents χ and $\bar{\chi}$.

Goods: One utility good $\hat{\gamma}_i$ for each agent i.

One deactivator good for a circle-of-death, $\hat{d}$.

Utilities: Agent i desires his utility good, i.e. $u_i(x) = x_{\hat{\gamma}_i}$.

Production: A circle of death $COD(\hat{d})$.

Director A wants Φ to be true and director B wants it to be false. The circle of death will ensure that Φ can never actually be falsified in a core allocation. (Note that the χ agents are defined only for the variables in x_1.)

In the first stage, director B decides whether to pick x_1 himself or defer to director A:

Goods: Authorization goods $\hat{a}_A$ and $\hat{a}_B$.

A choice source good $\hat{s}_{D,\chi}$ for each director D and variable χ in Φ.

Production: Director B has

$$Choice((2n_\Phi + 2) \cdot \hat{a}_A, (2n_\Phi + 1) \cdot \hat{a}_B)$$

Director A has

$$x_{\hat{d}} = f_{\hat{d}}(x) = x_{a_A}$$

$$Limit(\hat{d}, 1)$$

For each variable χ, each director D has

$$x_{s_{D,\chi}} = f_{s_{D,\chi}}(x) = x_{a_D}$$

$$Limit(\hat{s}_{D,\chi}, 2)$$

The director D whose good $\hat{a}_D$ is chosen will produce exactly two units of $\hat{s}_{D,\chi}$ for each variable χ. If A is chosen, he will also produce one unit of $\hat{d}$. (Because of the limit gadgets, the directors have no other way to consume all the authorization goods, so, following Proposition 2.4, we may assume that they do. This logic carries through the remainder of the construction.)

Next, the chosen director picks x_1. Both directors D have the following infrastructure (note that we use generalized choice gadgets as defined in Appendix A):

Goods: For each variable χ in x_1, two choice authorization goods $\hat{a}_{D,\chi}$ and $\hat{a}_{D,\bar{\chi}}$, and two assignment goods $\hat{\chi}$ and $\hat{\bar{\chi}}$.

For each variable χ in x_2, a choice source good $\hat{s}_{D,\chi}$ and two assignment goods $\hat{\chi}$ and $\hat{\bar{\chi}}$.

Production: Each director D owns, for each variable $\chi \in x_1$:

$$GChoice(\hat{s}_{D,\chi}, 3\hat{a}_{D,\chi}, 3\hat{a}_{D,\bar{\chi}})$$

and for each variable $\chi \in x_2$:

$$GChoice(\hat{s}_{D,\chi}, \hat{\chi}, \hat{\bar{\chi}}) \ .$$

Each χ agent has (for both directors):

$$x_{\hat{\chi}} = f_{\hat{\chi}}(x) = x_{\hat{a}_{D,\chi}}$$

$$Limit(x_{\hat{\chi}}, 1)$$

(the $\bar{\chi}$ agents have similar functions).

In essence, the director produces x_2 himself; however, though he chooses x_1, he must delegate the production of the $\hat{\chi}$ goods for x_1 to the χ and $\bar{\chi}$ agents. The extra authorization good $\hat{a}$ will later be used to generate utility.

Next Φ is evaluated:

Goods: True and false goods $\hat{\Phi}_{TRUE}$ and $\hat{\Phi}_{FALSE}$ to represent the value of Φ.

Production: Director B has a $SAT_\Phi(\{\hat{\chi}_i\}, \{\hat{\bar{\chi}}_i\}, \hat{\Phi}_{TRUE}, \hat{\Phi}_{FALSE})$ gadget to evaluate Φ given the $\hat{\chi}$ and $\hat{\bar{\chi}}$ goods.

Note that a coalition of B and the χ agents may evaluate Φ without the help of any other agents. Moreover, those same agents are required to compute Φ even if A is involved. This will restrict the possible defecting coalitions.

Finally, agents receive their payoffs:

Goods: For each χ (resp. $\bar{\chi}$) agent and each director D, an intermediate utility good $\hat{\alpha}_{D,\chi}$ (resp. $\hat{\alpha}_{D,\bar{\chi}}$).

Production: Director A has

$$x_{D,\hat{\gamma}_A} = f_{\gamma_A}(x) = \min((2n+2) \cdot x_{\Phi_{TRUE}}, x_{\hat{a}_A})$$

$$Limit(\hat{\gamma}_A, 1) \ .$$

Similarly, director B has

$$x_{\hat{\gamma}_A} = f_{\gamma_A}(x) = \min((2n+2) \cdot x_{\Phi_{FALSE}}, x_{\hat{a}_B})$$

$$Limit(\hat{\gamma}_A, 1) \ .$$

Each χ agent has, for each director D:

$$x_{\hat{\alpha}_{D,\chi}} = f_{D,\alpha_\chi}(x) = x_{\hat{a}_{D,\chi}}$$

$$Limit(\hat{\alpha}_{D,\chi}, 1) \ ,$$

(and $\bar{\chi}$ have similar functions).

Finally, each χ agent has

$$x_{\hat{\gamma}_\chi} = f_{A,\gamma_\chi,\chi}(x) = \min(x_{\hat{\alpha}_{A,\chi}}, (2n+2) \cdot x_{\Phi_{TRUE}})$$
$$x_{\hat{\gamma}_\chi} = f_{A,\gamma_\chi,\bar{\chi}}(x) = 2 \cdot \min(x_{\hat{\alpha}_{A,\bar{\chi}}}, (2n+2) \cdot x_{\Phi_{TRUE}})$$
$$x_{\hat{\gamma}_\chi} = f_{B,\gamma_\chi,\chi}(x) = \frac{3}{2} \min(x_{\hat{\alpha}_{B,\chi}}, (2n+2) \cdot x_{\Phi_{FALSE}})$$
$$x_{\hat{\gamma}_\chi} = f_{B,\gamma_\chi,\bar{\chi}}(x) = \frac{3}{2} \min(x_{\hat{\alpha}_{B,\bar{\chi}}}, (2n+2) \cdot x_{\Phi_{FALSE}})$$

(similar for the $\bar{\chi}$ agents).

Payoffs are summarized in the following table (D represents the chosen director):

Agent	$\Phi = 1$, $D = A$	$\Phi = 0$, $D = A$	$\Phi = 1$, $D = B$	$\Phi = 0$, $D = B$
A	1	0	0	0
B	0	0	0	1
χ produces x_1	1	0	0	$\frac{3}{2}$
χ does not produce x_1	2	0	0	$\frac{3}{2}$

Agents A and B fight over whether Φ is true or false. In allocations where A makes Φ true, the χ agents who don't participate in producing x_1 will get 2 units of utility and be happy, but the agents who do produce x_1 will only get 1 unit. Thus, if they can falsify Φ with B, the coalition would rather defect and get $\frac{3}{2}$. Note that if they defect, the only setting of x_1 that they can produce is the one originally chosen by A (because any other setting requires the production firms of the χ agents who got 2 and, therefore, would not want to join the coalition).

Since nobody gets any utility if Φ is not evaluated, we may assume that it is evaluated in any core allocation. Moreover, for the same reason, A must be the director if Φ is evaluated to true, and B must be the director if Φ is false. (A coalition of agents would clearly like to defect from an allocation in which nobody receives any utility.)

For the moment, let us assume that the Σ_2SAT instance is true, i.e. a "solution" x_1 exists. In this case, a core allocation is the one in which A picks an x_1 that solves the Σ_2SAT problem. In this case, the only agents who could reasonably defect are B and the χ agents who produce x_1. However, as noted, they are bound by A's choice of x_1, therefore they will not be able to falsify Φ and would receive no utility if they defect. Thus, the allocation is in the core. In contrast, if A picks the wrong x_1 and still tries to satisfy Φ, that same coalition will defect and falsify Φ on its own. Finally, any allocation in which A is not in charge is precluded by the circle-of-death. The case that no solution x_1 exists is merely a subset of the cases mentioned above.

Thus, any core allocation must correspond to an x_1 that "solves" the Σ_2SAT instance. It follows that core allocations are $F\Sigma_2^P$-complete to compute.

□

Proof: (of Theorem 3.3) Given a SAT formula Φ, we construct an economy with the following properties: if Φ is unsatisfiable, the economy has a trivial price equilibrium. If Φ is satisfiable, then we get the economy in Section 4 of [6] that has no price equilibrium. The main trick is to manipulate the set of production possibilities for the two goods $\hat{a}$ and $\hat{b}$: when Φ is satisfiable, we want it to be $T = \{x_{\hat{a}}, x_{\hat{b}} | x_{\hat{a}} \leq 2 \text{ and } x_{\hat{b}} \leq 2 \text{ and } (x_{\hat{a}} \leq 1 \text{ or } x_{\hat{b}} \leq 1)\}$ (see the picture in [6]), and when it is unsatisfiable, we want it to be $F = \{x_{\hat{a}}, x_{\hat{b}} | x_{\hat{a}} \leq 1 \text{ and } x_{\hat{b}} \leq 1\}$ (this is a convex set, so there will be a price equilibrium).

To accomplish this, we construct a SAT gadget, as above. When Φ is false, we have a false good. We give agents the technology to turn this good into any quantity of goods $\hat{a}$ and $\hat{b}$ from set F.

Since this set is convex (and the other goods, i.e. those used in the SAT formula are not traded), the economy has a price equilibrium. Now, when Φ is true, we allow the economy to produce from T as follows: a choice gadget produces one of two intermediate goods. The first can be used in production of any vector of goods up to $\hat{a} = 1$ and $\hat{b} = 2$, while the other can be used in production of any vector of goods up to $\hat{a} = 2$ and $\hat{b} = 1$. It is straightforward to integrate the utilities and endowments to make the example work.

□

B.2 Sequential Production with Isolation

Proof: *(of Theorem 5.2.)* Membership in $\mathrm{F}\Delta_3^P$ is the easy direction. As with core defections, given a plan that represents a defection, it is easy (efficiently computable) to check. It follows that an individually-rational production plan may be found in NP with an NP oracle (guess the plan and use the oracle to verify its rationality). To find the optimal such plan, we maximize social welfare. Thus, the problem lies in a generalization of OptP to the second level of the polynomial hierarchy and, therefore, in $\mathrm{F}\Delta_3^P$.

To prove that this problem is $\mathrm{F}\Delta_3^P$-hard, we reduce from the $\mathrm{F}\Delta_3^P$-complete problem lexicographically maximum Σ_2SAT [15], i.e. find the lexicographically maximum x_1 such that for all x_2, $\Phi(x_1, x_2) = 1$ (where Φ has n_Φ variables). We define the following economy:

Agents: Two agents: agent A and B.

Goods: Two utility goods $\hat{\gamma}_A$ and $\hat{\gamma}_B$.

An initial seed good $\hat{s}_A$.

Approval goods $\hat{a}$ and $\hat{a}'$.

For each $\chi \in \Phi$, a source good $\hat{s}_\chi$.

For each $\chi \in \Phi$, assignment goods $\hat{\chi}$ and $\hat{\bar{\chi}}$.

True and false goods $\hat{\Phi}_{TRUE}$ and $\hat{\Phi}_{FALSE}$.

For each agent and value of Φ, intermediate utility goods $\hat{\alpha}$, e.g. $\hat{\alpha}_{A,TRUE}$.

For each $\chi \in x_1$, an "intermediate counting good" $\hat{\beta}_\chi$.

Utilities: $u_i(x) = x_{\hat{\gamma}_i}$.

Endowments: For each $\chi \in x_1$, $e_{A,\hat{s}_\chi} = 2$.

Production: For each $\chi \in x_1$, agent A has

$$GChoice(\hat{s}_\chi, [2\hat{\chi}, \hat{a}], [2\hat{\bar{\chi}}, \hat{a}], \frac{1}{n_\Phi}\hat{\gamma}_A) \ .$$

Agent B has

$$x_{\hat{a}'} = f_{a'}(x) = |x_2| \cdot \max(x_{\hat{a}} - (|x_1| - 1), 0)$$

and for each $\chi \in x_2$

$$x_{\hat{s}_\chi} = f_{s_\chi}(x) = 2 \cdot x_{\hat{a}'}$$

$$GChoice(\hat{s}_\chi, \hat{\chi}, \hat{\bar{\chi}}) \ .$$

Agent B also has $SAT_\Phi(\{\hat{\chi}_i\}, \{\hat{\bar{\chi}}_i\}, \hat{\Phi}_{TRUE}, \hat{\Phi}_{FALSE})$ to evaluate Φ.

For payoffs, agent A has

$$x_{\hat{\alpha}_{A,TRUE}} = f_{\alpha_{A,T}}(x) = 4 \cdot x_{\hat{\Phi}_{TRUE}}$$

$$Limit(\hat{\alpha}_{A,TRUE}, 1)$$
$$x_{\hat{\gamma}_B} = f_{\gamma_b}(x) = 2 \cdot x_{\hat{\alpha}_{A,TRUE}}$$

and agent B has

$$x_{\hat{\alpha}_{B,FALSE}} = f_{\alpha_{B,F}}(x) = 4 \cdot x_{\hat{\Phi}_{FALSE}}$$
$$Limit(\hat{\alpha}_{B,FALSE}, 1)$$
$$x_{\hat{\gamma}_B} = f_{\gamma_b}(x) = 2 \cdot x_{\hat{\alpha}_{B,FALSE}} \ .$$

Meanwhile, agent A has the following for each variable $\chi \in x_1$:

$$x_{\hat{\beta}_\chi} = f_{\beta_\chi}(x) = \min(x_{\hat{\chi}}, 2n_\Phi \cdot \Phi_{TRUE})$$

$$Limit(\hat{\beta}_\chi, 1) \ .$$

and a binary counting gadget

$$BCG(\hat{\gamma}_A, \hat{\beta}_{\chi_0}, \hat{\beta}_{\chi_1}, \dots) \ .$$

The economy functions in three stages. First, A picks x_1. Once x_1 is chosen, B picks x_2 (the structure of approval goods $\hat{a}$ ensures that no variable in x_2 is chosen before all variables in x_1 have been fixed). Finally, payoffs are computed based on the results of evaluating Φ. Agent B receives 2 units if it is false, and agent A receives $2 + \sigma$ units if Φ is true, where σ is the value found by taking x_1 as the binary representation of a number. Agent A also has a default option to refuse to produce x_1, thereby generating a small amount of utility $0 < u_A < 1$. (The complicated structure of intermediate goods merely ensures the correct discretization and distribution of goods.)

Production plans come in three flavors: satisfy Φ, falsify Φ, or neither. If Φ is satisfied, then agent A will be happy. However, agent B would rather falsify Φ. Thus, if B can pick x_2 so that Φ is false, he will defect, since he does not need to interact with any other players once he has the goods specifying x_1. The only individually rational plan that satisfies Φ will include an x_1 such that $\Phi(x_1, x_2)$ is always true.

In plans that falsify Φ, agent A gets nothing and, therefore, will defect at the beginning and choose the default option. Thus, no such plan may be individually rational.

Finally, in plans that do not compute Φ, agent B does not receive any utility, and A receives at most 1 unit (from the default option). This is strictly dominated by any plan in which Φ is made to true, so it can only be the optimal individually rational plan if for any x_1, there exists an x_2 such that $\Phi(x_1, x_2)$ is false.

Thus, if there is an x_1 such that $\Phi(x_1, x_2)$ is always true, the optimal individually rational production plan is the one in which A picks the x_1 that maximizes σ, i.e. the lexicographically maximum x_1. In other cases, the only individually rational option is for A to take the default option, signaling that no such x_1 exists. It follows that computing an optimal individually rational sequential production plan is equivalent to lexicographically maximum Σ_2SAT and, therefore, is $F\Delta_3^P$-complete.

□

C Average-Case Complexity Equilibrium

Proof: (Rough sketch for Theorem 6.4.) Consider modeling an economy as a set of agents on a directed graph $G = (V, E)$. Each vertex $v \in V$ corresponds to an agent who, at time t, takes input from incoming edges and distributes output on outgoing edges. Each agent may also elect to keep some goods for himself (these are treated as extra inputs at time $t+1$). The economy is in a feasible configuration if the quantity of goods taken as input by vertex v equals the quantity of goods sent to v as output by neighbors of v at time t. As in the choice gadget, we can use non-convexity to force each agent to pick a discrete production task with his inputs.

In the Arrow-Debreu framework, we model this economy by creating a separate copy of the goods and production functions for each time t and vertex v. (Note that this is precisely how Arrow and Debreu [3] suggest modeling an economy over time and space.)

Hardness will follow because this economic model is a superset of an edge tiling problem defined by Gurevich [10]. An edge tiling problem consists of a set of tiles T, an $n \times n$ square, and some initial conditions. The goal is to place one tile at each location in the square such that adjacent labels match and all initial constraints are satisfied.

Gurevich [10] shows that when the first row is randomly filled according to a certain "uniform" distribution (i.e. the initial conditions) and all possible sets of tiles T occur with positive probability, it is average-case NP-complete to decide if the $n \times n$ square may be tiled.

The edge tiling problem corresponds to an economy where agents are organized on a line. Each agent chooses his production task (i.e. his tile) such that his inputs and outputs (i.e. the labels on the left and right sides of a tile) match his neighbors. Each row of the tiling represents a time t, so labels on the top and bottom edges of the tile correspond to goods saved by an agent. Similar to the construction in Theorem 6.3, players only receive a payoff if the square is completely tiled.

Reducing from Gurevich shows that when production at time $t = 0$ (i.e. filling the first row) is done according to the proper distribution, *improving from payoff 0 in such an economy is average-case NP-hard for any distribution over possible production tasks such that every set of tiles occurs with positive probability.*

The exact construction leverages fixed costs as in choice gadget to force discrete choices. □

The economy in this proof is somewhat reasonable. It is powerful because the precise distribution does not matter, provided it satisfies the very general condition that all sets of tiles are possible. The linearity requirement is somewhat suspect, but is more an artifact of this particular proof than a fundamental requirement for average-case hardness.

The main drawback of this construction is that it requires randomization over discrete production tasks. In contrast, a more natural form of randomization would fix the discrete choices and randomize over continuous parameters of those choices. We do not know of a natural construction that does this.

References

[1] C. D. Aliprantis, B. Cornet, and R. Tourky. Economic equilibrium: Optimality and price decentralization. *Positivity*, 6(3):205–241, 2002.

[2] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach.* Cambridge University Press, New York, NY, USA, 2009.

[3] Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.

[4] Andrej Bogdanov and Luca Trevisan. Average-case complexity. *CoRR*, abs/cs/0606037, 2008.

[5] J.-M. Bonnisseau and B. Cornet. Existence of equilibria with a tight marginal pricing rule. *Journal of Mathematical Economics*, 44(7-8):613 – 624, 2008. Special Issue in Economic Theory in honor of Charalambos D. Aliprantis.

[6] Donald J. Brown. Chapter 36 equilibrium analysis with non-convex technologies. volume 4 of *Handbook of Mathematical Economics*, pages 1963 – 1995. Elsevier, 1991.

[7] Bruno Codenotti, Amin Saberi, Kasturi Varadarajan, and Yinyu Ye. Leontief economies encode nonzero sum two-player games. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 659–667, New York, NY, USA, 2006. ACM.

[8] Luis C. Corchn and M. Socorro Puy. Individual rationality and voting in cooperative production. *Economics Letters*, 59(1):83 – 90, 1998.

[9] Roger Guesnerie. Pareto optimality in non-convex economies. *Econometrica*, 43(1):1–29, 1975.

[10] Yuri Gurevich. Average case completeness. *JOURNAL OF COMPUTER AND SYSTEM SCIENCES*, 42:346–398, 1991.

[11] Kamal Jain and Kasturi Varadarajan. Equilibria for economies with production: constant-returns technologies and production planning constraints. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 688–697, New York, NY, USA, 2006. ACM.

[12] Kamal Jain, Vijay V. Vazirani, and Yinyu Ye. Market equilibria for homothetic, quasi-concave utilities and economies of scale in production. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 63–71, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[13] Sanjiv Kapoor, Aranyak Mehta, and Vijay Vazirani. An auction-based market equilibrium algorithm for a production model. *Theor. Comput. Sci.*, 378(2):153–164, 2007.

[14] M W Krentel. The complexity of optimization problems. In *STOC '86: Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 69–76, New York, NY, USA, 1986. ACM.

[15] Mark W. Krentel. Generalizations of opt p to the polynomial hierarchy. *Theor. Comput. Sci.*, 97(2):183–198, 1992.

[16] David M. Kreps. *A Course in Microeconomic Theory.* Princeton University Press, February 1990.

[17] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory.* Oxford University Press, June 1995.

[18] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, New York, NY, USA, 2007.

[19] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.

[20] Grant Schoenebeck and Salil Vadhan. The computational complexity of nash equilibria in concisely represented games. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 270–279, New York, NY, USA, 2006. ACM.

[21] Guoqiang Tian. On the informational requirements of decentralized pareto-satisfactory mechanisms in economies with increasing returns. October 2006.

[22] K. Vela Velupillai. Uncomputability and undecidability in economic theory. *Applied Mathematics and Computation*, In Press, Corrected Proof:–, 2009.