Electronic Colloquium on Computational Complexity, Report No. 17 (2007)

Indivisible Markets with Good Approximate Equilibrium Prices *

Richard Cole
Computer Science Department
Courant Institute of Mathematical Sciences
New York University

Ashish Rastogi
Computer Science Department
Courant Institute of Mathematical Sciences
New York University

Abstract

This paper considers the tradeoff between divisibility and the hardness of approximating equilibrium prices. Tight bounds are obtained for *smooth* Fisher markets that obey a relaxed weak gross substitutes property (WGS). A smooth market is one in which small changes in prices cause only proportionately small changes in demand, which we capture by a parameter k. Specifically, assuming that the total wealth is at least r times the total number of goods, this paper gives a polynomial time algorithm to compute prices achieving a $(1 - O(k/r))$-approximation and shows that it is NP-hard to do better.

A second contribution of this paper is a new consideration of how to measure the quality of an approximation to equilibrium prices. Our approach takes the notion of compensatory payments from welfare economics and applies it to indivisible markets. This allows the dissatisfaction, or *discontent*, of individual agents to be combined in a natural way. In addition, an important observation is that in the indivisible setting, standard utility functions, such as CES, need not obey the standard WGS property.

*This work was supported in part by NSF grant CCF 0515127.

ISSN 1433-8092

1 Introduction

Deng, Papadimitriou and Safra [8] showed that it is APX-Hard to approximate equilibrium prices and allocations for indivisible Arrow-Debreu markets. Really, they show that the optimal allocation is hard to approximate. Also, although they do not state this, their construction proves the result for the more restricted Fisher market. Subsequently, there has been a considerable body of work giving polynomial time algorithms to compute approximate equilibria for divisible markets [8, 6, 4, 13, 10]. Of course, all markets are necessarily indivisible. While one anticipates large numbers of copies of each good, and prices with a minimum adjustment that is a small fraction of the price, nonetheless both these quantities are bounded. In our view, an implicit assumption in studying the divisible problem is that the discreteness present in actual markets has only a small effect. As in practice the number of copies of a good need not be much greater than the number of interested buyers, it is not clear that the assumption is sound in general. Nonetheless, this paper does show that near-optimal prices can be found efficiently in the markets we consider.

Deng et al. [8] also considered algorithms for indivisible markets. They gave an exhaustive algorithm for computing an approximate equilibrium in polynomial time for markets with a constant number of distinct goods. This paper gives a simple efficient algorithm for indivisible markets with no constraints on the number of goods or buyers.

In Arrow-Debreu markets each agent is provided with an initial allocation of goods. The problem is to find prices at which every agent can trade its initial allocation for an optimal bundle in such a way that supply and demand exactly balance. In a divisible setting, where all goods are infinitely divisible and prices can take on any real value, under modest assumptions, equilibrium prices always exist [2]. This need not be the case in an indivisible setting. The task then becomes one of finding prices that support a near-equilibrium allocation.

In this paper, we limit ourselves to Fisher markets. In these markets, there are two groups of agents: buyers, who initially have only money, and sellers, who have only goods. Each seller's goal is to sell all its goods. Each buyer's goal is to utilize its money to obtain an optimal bundle of goods, which may include money (this is a generalization of the standard model where buyers have no desire for money). In fact, we use a small variant of the Fisher model in which buyers may attach utility to money, and thus an optimal allocation may leave them with non-zero money at the end. This strikes us as reflecting actual behavior. Devanur and Vazirani [9] had been concerned with a related issue, namely that in a Fisher market with linear utilities, buyers might spend all their money on one good. They proposed putting limits on spending and also incorporating a desire for money in the context of utility functions that are linear step functions. We note that our results still hold if a buyer has zero utility for money.

Let $\mathbf{p} = (p_1, \ldots, p_n)$ be a collection of prices. Whenever convenient, we write $\mathbf{p} = (p_i, \mathbf{p}_{-i})$, where $\mathbf{p}_{-i}$ is the collection of prices for all goods except good i. Further, we let $x_i(\mathbf{p})$ denote the demand for good i at $\mathbf{p}$.

We limit the markets we consider to obey a discrete version of the following bound: the (fractional) rate of change of spending on a good with respect to a (fractional) change in its price is bounded by a non-positive constant. Without such a condition it is not clear how to design algorithms for finding approximate equilibria that change prices by (multiplicative) increments. One way of expressing this bound in the divisible setting is that there is a constant $l \geq 0$ such that:

$$p_i x_i(p_i, \mathbf{p}_{-i}) \leq (p_i + h) x_i(p_i + h, \mathbf{p}_{-i}) \left(1 + h/p_i\right)^l \text{ for all } h > 0$$

or equivalently, there is a constant $k \geq 1$ such that

$$x_i(p_i, \mathbf{p}_{-i}) \leq x_i(p_i + h, \mathbf{p}_{-i})\,(1 + h/p_i)^k \text{ for all } h > 0.^1 \tag{1}$$

Also, as $h \to 0$, the first expression can be rewritten as:

$$\frac{(p_i + h)x_i(p_i + h, \mathbf{p}_{-i}) - p_i x_i(p_i, \mathbf{p}_{-i})}{(p_i + h)x_i(p_i + h, \mathbf{p}_{-i})} \Big/ \frac{h}{p_i} \geq -l,$$

which is the desired derivative.

k provides a bound on the elasticity of demand for $x_i(\cdot)$ with respect to p_i[15]. We now modify this definition for the discrete setting to avoid fractional demands, and call it k-smoothness.

Definition 1. *A market is k-smooth if for every good i, for all prices $\mathbf{p}$,*

$$x_i(p_i, \mathbf{p}_{-i}) \leq \lceil x_i(p_i + h, \mathbf{p}_{-i})\,(1 + h/p_i)^k \rceil, \textit{ for all } h > 0 \textit{ such that } x_i(p_i + h, \mathbf{p}_{-i}) \geq 1,$$

and $x_i(p_i, \mathbf{p}_{-i}) \leq 1$ if $x_i(p_i + 1, \mathbf{p}_{-i}) = 0$.

Note that $x_i(p_i, \mathbf{p}_{-i}) \leq 2^k x_i(2p_i, \mathbf{p}_{-i})$ if $x_i(2p_i, \mathbf{p}_{-i}) \geq 1$.

Summary of results. Our main results are:

- A polynomial time algorithm that finds prices for a k-smooth market such that the resulting allocation of goods is a $1 - O(k/r)$ approximation (in a sense made precise later). This algorithm assumes a polynomial time oracle which, given the prices for every good, returns the demand for each good. The algorithm makes $O(ktn \log w)$ calls to the oracle, where n is the number of distinct goods in the market, w is the total buyer wealth, and t is the total number of copies of goods in the market; the oracle calls dominate the overall running time. (If $t \gg r$, the running time can be made linear in r rather than t. The running time can also be reduced proportionately if a less good approximation is sought). An appealing aspect of the algorithm is its simplicity.

- It is NP-hard to find prices at which a $1 - \Theta(k/r)$-approximation is possible. The hardness lies in finding the prices, not the allocation. To the best of our knowledge, this is the first result demonstrating that it is hard to find correct prices as opposed to a correct allocation. We also remark that in the present setting it seems much more delicate to obtain the hardness result for prices.

Our algorithm assumes a model in which each buyer seeks certain goods and either receives them or not depending on the supply. There is no substitution of second best choices. On the other hand, our hardness result holds even if arbitrary allocations are permitted. Since the results only make sense for $k/r \leq 1$, we assume $k \leq r$ for the rest of this paper.

k-smoothness in the indivisible setting corresponds in the divisible setting to bounding the rate of change of demand for a good with respect to its price. Assuming that the demand function

[1] Note the internal consistency of this definition; that is,

$$(1 + (h_1 + h_2)/p_i) = (1 + h_1/p_i)\,(1 + h_2/(p_i + h_1)).$$

for every good has a bounded partial derivative with respect to every price, Codenotti et al. [6] gave strongly polynomial time algorithms for finding approximate equilibria using a tatonnement-type procedure. The algorithm of Codenotti et al. uses the ellipsoid method as a subroutine; this contrasts with the simplicity of our method. Of course the algorithms are not directly comparable as they are for different settings, and as we discuss below, use different approximation measures. Another point to keep in mind is that a rounded optimal or near-optimal solution for the divisible setting need not be a near-optimal solution in the indivisible setting (assuming analogous problems can be defined). This can occur, for example, when there are goods for which buyers have $o(1)$ demand, even if there are many copies of the good.

k-smoothness constrains the behavior of the overall demand in the market. Without some constraints, equilibria need not be unique and can be hard to compute [7]. Most previous algorithmic work in the divisible setting overcame this by constraining individual utility functions (e.g. to satisfy the weak gross substitutes property, or to be linear [6, 4]). However, this can still result in some quite unintuitive behavior. For example, if all the buyers had identical linear utility functions, a minimal change in prices can shift the demand from being all for one good to all for another good. k-smoothness allows linear utilities, but it limits the amount of a pair of goods that are equally desired at a given set of prices (thereby avoiding large fluctuations in demand with small changes in price), and so in particular precludes all buyers having identical linear utilities.

Indivisible markets have been studied by mathematical economists also. Ausubel, Gul and Stacchetti [3, 11] restrict the utility functions being allowed so as to ensure that equilibria exist. In fact, they suppose the buyers have valuation functions that are integer valued, i.e. each basket of goods has an associated dollar value. They show that if these functions also satisfy an "individual substitutes" property, then there is a Walrasian equilibrium. The individual substitutes condition is a further restriction of the well known gross substitutes condition. Recall that the gross substitutes condition requires that if the prices of some goods are raised while the prices of the remaining goods are held constant, then an agent's demand increases weakly for each of the goods whose prices were held constant. The individual substitutes property in addition requires separate prices for distinct copies of the same good and that individual copies of the same good also obey the gross substitutes property. Under these restrictions on the utility functions, they show that the minima of a certain potential function (a Lyapunov function) correspond to equilibrium prices. They propose a tatonnement algorithm that updates prices at discrete time steps in the direction of decreasing potential function values. The algorithm exhaustively evaluates the potential function at every price point in a unit neighborhood of the current price. Integer valuations imply that the potential function decreases by an integer amount at each step. Thus, the algorithm converges to equilibrium prices, but does so in exponential time.

In [16] Milgrom and Strulovici consider the above setting, replacing the individual substitutes constraint with the standard WGS constraint. As an equilibrium may then not exist they propose a notion of pseudo-equilibrium. The authors argue that a pseudo-equilibrium price is also an *approximate* equilibrium price by showing that the excess demand (which should be zero at equilibrium prices) can be bounded by a function of the number of goods, the number of buyers and the largest *gap* in demand for each good among optimal bundles at any price. They also give an exhaustive algorithm. This approach is in contrast with our work which seeks to estimate individual discontent in terms of money and aggregate the overall discontent in the unit of money.

In the divisible setting, it is well known that the market obeys the weak gross substitutes property when buyers in the market possess linear, Cobb-Douglas or CES utility functions with

$0 < \rho < 1$. A CES utility function orders a buyer's preferences over bundles of goods $\mathbf{x}$ according to the value $u(\mathbf{x}) = (\sum_i \alpha_i x_i^{\rho})^{1/\rho}$, with $-\infty < \rho < 1, \rho \neq 0$. The linear utility function and the Cobb-Douglas utility function arise as special cases of the CES utility function when $\rho = 1$ and $\rho \to 0$ respectively. However the WGS property need not hold in the indivisible setting and the indivisibility can result in quite unintuitive behavior, as we illustrate later via an example. To account for such effects, we consider a relaxation of the WGS property and our algorithm and the hardness result assume that the overall market obeys the relaxed WGS property.

Next, we discuss how to measure the quality of an approximate equilibrium. Previous work had measured it by considering each agent in turn and asking what is the ratio between its actual utility and the maximum utility it could conceivably achieve at the prices on offer; the quality of the approximation is then defined to be the minimum of these ratios over all agents (0 is the worst possible, 1 the best). We offer two critiques of this approach.

The first is that utilities are being treated implicitly as if they were valuation functions. But this is ascribing more meaning to the function than may be appropriate, since in general there are multiple utility functions that represent a given well behaved preference ordering. Indeed, if u is a utility function and f is a strictly increasing function, then $f(u)$ provides an alternative utility function giving the same preference ordering.[2] This is disconcerting, for a $(1-\epsilon)$-approximation under one utility function may be a much better or worse quality approximation for another equivalent utility function.

Our second critique concerns how the utilities are combined. In our view, the quality of the approximation ought to correspond to how dissatisfied the market participants are collectively; surely participants with large resources will have a larger impact. This suggests an approach other than minimizing relative discontent. Instead, we seek to combine the discontent of each of the participants. But this necessitates expressing discontent in a common unit. We do this using the notion of compensatory payments from welfare economics: essentially, this asks what is the difference in value to the agent between the agent's optimal allocation at the current prices and the allocation the agent receives. We call this quantity, measured in the unit of money, the agent's *discontent*.

It is natural to ask what is the relationship between an individual's discontent in our measure and that given by comparing utility values. We show that in the divisible setting, for concave utility functions, for a given allocation, the approximation factor is always at least as large in our measure, and can be arbitrarily worse. In other words, for each individual, a given approximation factor is at least as hard to achieve in our measure as in the utility measure.

Once one has a way to measure the individual discontent of each participant, the next issue is to aggregate individual discontent to represent the overall discontent of the market. Here, the fact that our measure of individual discontent is in the common unit of money allows us to simply combine discontents additively.

This can be viewed as an L_1 norm; our hardness results immediately carry over to any L_p norm, $p > 1$, and to the L_∞ or maximum norm. Loosely speaking, the L_∞ norm is analogous to but not the same as the previous approximation measure. Our algorithmic results do not immediately extend to these other norms.

[2]When studying choice under risk (for example in auctions), the concavity of u defines agents' risk aversion. In such cases, the utility function may be unique. However, this is not the case in a market setting, where we concern ourselves with deterministic outcomes.

We conclude the introduction by summarizing our contributions. An important part of the paper lies in the definitions it introduces: k-smoothness and discontent. In our view, k-smoothness provides a natural way to specify well-behaved markets, and the definition of discontent supports a more robust notion of approximation of equilibria than the previous approach. The paper also gives a simple polynomial time algorithm for computing approximate equilibria and a complementary NP-hardness result.

In Section 2, we provide formal definitions. Section 3 gives our algorithm and Section 4 the hardness results.

2 Definitions

Definition 2. (Utility Function.) *Let G denote the set of goods present in the market. Then $u : \mathbb{R}^G \to \mathbb{R}^+$ is said to be a utility function, if, for $A_1, A_2 \subseteq G$, the allocation A_1 is strictly preferred to A_2 exactly when $u(A_1) > u(A_2)$.*

Comment 1. *In the auction literature, utility is often defined to be the difference between valuation and price; this is* not *the meaning intended here.*

Definition 3. (The Market.) *The market consists of a set G of n goods in supply $\mathbf{s}$ (s_g copies of good g for $g \in G$), a set B of m buyers, buyer wealth $\mathbf{w}$ (an initial amount w_b for $b \in B$). Each buyer b possesses a utility function u_b over the basket of goods. $(G, \mathbf{s}, B, \mathbf{w}, (u_1, \ldots, u_m))$ denotes an instance of the market.*

Comment 2. *In addition to the indivisibility of goods, we also impose indivisibility of money in the market. Thus, money is no longer* fluid*, and prices can be set only at indivisible integer values. Henceforth, for specificity, we take the unit of money to be a dollar.*

Given a market M, the problem is to find prices, called market-clearing prices, such that at these prices each buyer receives its optimal allocation and no goods are left unsold (i.e. all sellers are also optimally happy). Formally, this can be written as:

Indivisible extended Fisher market problem. Given a market $(G, \mathbf{s}, B, \mathbf{w}, (u_1, \ldots, u_m))$, determine prices $\mathbf{p}$ such that there exists a partitioning of goods to buyers, with buyer b receiving A_b, $A_b \cdot \mathbf{p} \leq w_b$, $\bigcup_b A_b = G$, and for each b, A_b maximizes b's utility: for all A'_b such that $A'_b \cdot \mathbf{p} \leq w_b, u_b(A_b, w_b - A_b \cdot \mathbf{p}) \geq u_b(A'_b, w_b - A'_b \cdot \mathbf{p})$.[3]

Given a set of prices, for each buyer there is a basket of goods which maximizes its utility. An optimal allocation is one which maximizes every buyer's utility simultaneously and leaves no goods unsold. In the indivisible setting, in general, there need not be any price collection at which an optimal allocation exists, which raises the question of how far from optimal a given allocation is. To this end, we define a notion of individual discontent.

Our approach is based on the notion of compensatory payments from welfare economics. It asks what payment an agent a needs to receive to compensate for a non-optimal allocation. Previously, this was defined for the divisible setting, as follows. Let A^{opt} be an optimal allocation, A^{act} the actual allocation, and A^l a least cost allocation with $u(A^l) = u(A^{act})$. Note that the choice of A^l depends on the prices and on agent a's utility function. Then the compensation agent a needs is

[3] In computing utility, we view money as just another good. However, to conform with the usual perspective, we list it separately from the other goods.

defined to be $\text{cost}(A^{opt}) - \text{cost}(A^l)$. We call $\text{cost}(A^l)$ the *value* of A^{act} for agent a at the current prices.

Our approach is analogous. For each agent a (buyer or seller), for each amount of money, \$1, \$2, …, \$$i$, …, we consider an optimal allocation A_a^i of goods and m_a^i of money, its utility $u_a(A_a^i, m_a^i)$ and its cost for agent a, $A_a^i \cdot \mathbf{p} + m_a^i = i$ at the given prices. It may be that $A_a^i = A_a^{i+1}$ and $u_a(A_a^i, m_a^i) = u_a(A_a^{i+1}, m_a^{i+1})$ for some i (this can occur only if agent a has no utility for money, and there is nothing useful to purchase with the last dollar). To find the *value* of an allocation A with money m_a (written as (A, m_a)) to agent a, we find the least i such that either $u_a(A_a^i, i - A_a^i \cdot \mathbf{p}) = u_a(A, m_a)$ or $u_a(A_a^i, i - A_a^i \cdot \mathbf{p}) < u_a(A, m_a) < u_a(A_a^{i+1}, i+1 - A_a^{i+1} \cdot \mathbf{p})$. In the former case, the value of (A, m_a) is \$$i$, whereas in the latter case, the value of (A, m_a) is defined to lie between \$$i$ and \$$i+1$. In the second case, given that money is indivisible, it does not appear possible to define the value more precisely; in this situation \$$i$ is called the *lower-value* and \$$i+1$ the *upper-value* of (A, m_a) at prices $\mathbf{p}$. When the agent is a buyer (say buyer b with wealth w_b), we let $v_b(A, w_b - A \cdot \mathbf{p})$ denote the *upper-value* of allocation A at prices $\mathbf{p}$, defined to be the upper value of allocation A together with leftover money $w_b - A \cdot \mathbf{p}$. Note that for the optimal allocation, the upper value is the same as the lower value.

Henceforth, in scenarios in which a buyer has no marginal utility for money (i.e. adding money to a buyer's allocation does not improve her utility), we use the term optimal allocation to mean a maximum utility allocation of least cost.

Definition 4 (Discontent). *Let $(A, w_a - A\cdot\mathbf{p})$ be an allocation and (A^{opt}, m^{opt}) an optimal allocation (both including money) to agent a with wealth w_a at prices $\mathbf{p}$. The* upper-discontent *of a with allocation $(A, w_a - A \cdot \mathbf{p})$ at prices $\mathbf{p}$, $ud_a((A, w_a - A \cdot \mathbf{p}), \mathbf{p})$, is given by $(A^{opt} \cdot \mathbf{p} + m^{opt}) -$ lower-value$_a(A, w_a - A \cdot \mathbf{p})$.* Lower-discontent, *$ld_a$, is defined analogously. We use the notation d_a for short when no ambiguity will result.*

We observe that the seller discontent has a very simple form. First, we assume, without loss of generality, that there is a single distinct seller for each distinct good. If there are c_g copies of good g unsold at price p_g, the seller v_g for good g has a discontent of $d_{v_g}(c_g, \mathbf{p}) = c_g p_g$.

Note that upper-discontent minus lower-discontent is either \$0 or \$1. If they are equal, they are both called *discontent* for short. Further, for buyer b with wealth w_b, both lower and upper discontent lie in the range $[0, v_b(A^{opt}, m^{opt})] = [0, v_b(A^{opt}, w_b - A^{opt} \cdot \mathbf{p})]$.[4] In our algorithms, we will use upper discontent as a worst case measure of discontent, while for the hardness results, we will use lower-discontent, as it provides a lower bound.

From an agent's point of view, how well it is doing, its *efficiency*, is the ratio of value achieved (lower-value say) to the value of its optimal allocation.

We define *market efficiency* as the ratio of the utility derived by all the agents compared to what appears possible individually at the offered prices. More precisely:

Definition 5. (Market Efficiency.) *The market efficiency of an allocation A under prices $\mathbf{p}$ is defined as:*

$$1 - \frac{\sum_{a \in \mathcal{A}} d_a((A_a, w_a - A_a \cdot \mathbf{p}), \mathbf{p})}{\sum_{a \in \mathcal{A}} v_a(A_a^{opt}, m_a^{opt})}$$

where $\mathcal{A}$ denote the set of agents in the markets (all buyers and sellers).

[4] $m^{opt} \leq w_b - A^{opt} \cdot \mathbf{p}$, with equality when b has a desire for every remaining dollar in the leftover money.

Another view of market efficiency is that it is simply a weighted average of the agent efficiencies, where an agent's weight is the value of its optimal allocation.

If all prices are zero, the market efficiency is not well defined (it is $1 - 0/0$). To avoid this difficulty, we limit the allowable prices to be strictly positive. Further, note that if the market has efficiency 1, then an equilibrium has been achieved, for the sellers will have sold all their goods and every buyer will have an optimal allocation. Conversely, if an equilibrium has been achieved, then the market has an efficiency of 1.

Indivisible Market Value Problem. Given a market $M = (G, \mathbf{s}, B, \mathbf{w}, (u_1, \ldots, u_m))$, determine prices $\mathbf{p}$ and an allocation of goods to buyers that maximizes the market efficiency.

Our results are parameterized by r, the average wealth per unit good. Formally, $r = \frac{\sum_{b \in B} w_b}{\sum_{g \in G} s_g}$.

3 The Algorithm

In this section, we present an algorithm that approximates the efficiency of a k-smooth market to a factor $1 - O(k/r)$. We assume that an oracle for computing the aggregate excess demand is available. This assumption is fairly standard [6, 12, 5, 14]. Let $\mathbf{p} = (p_1, \ldots, p_n)$. For ease of exposition, we first describe the algorithm assuming that the market obeys the WGS property. We then give our definition of the relaxed WGS property and describe a straightforward modification of the algorithm to operate in the latter setting.

Recall that a market is said to be k-smooth if for every good i, for all prices $\mathbf{p}$,

$$x_i(p_i, \mathbf{p}_{-i}) \leq \lceil x_i(p_i + h, \mathbf{p}_{-i}) \, (1 + h/p_i)^k \rceil, \text{ for all } h > 0 \text{ such that } x_i(p_i + h, \mathbf{p}_{-i}) \geq 1.$$

and $x_i(p_i, \mathbf{p}_{-i}) \leq 1$ if $x_i(p_i + 1, \mathbf{p}_{-i}) = 0$.

Definition 6. (High price.) *A price p_i in a collection $\mathbf{p} = (p_1, \ldots, p_n)$ of prices is said to be* high *if there is a surplus of good i at $\mathbf{p}$. The surplus need not be strict.* (Barely high price.) *Price p_i is said to be* barely high *if there is a surplus of good i at $\mathbf{p}$, but a deficit of good i at $(p_i - 1, \mathbf{p}_{-i})$.*

Low and barely low prices are defined analogously.

Let w denote the total buyer wealth and s be the minimum number of copies of each good present in the market; i.e. $w = \sum_i w_i$, $s = \min_j s_j$. Let $x_i(\mathbf{p})$ denote the demand for good i at prices $\mathbf{p}$. We aim to find a collection of prices $\mathbf{p}$ so that all goods are priced *barely high.* The algorithm in Figure 1 computes such prices while keeping all prices high at all times. On termination, as all the prices are barely high, the supply of each good is at least the demand, and thus each buyer can be and is allocated all the goods it seeks. We will show that this allocation achieves an efficiency of $1 - O(k/r)$.

Algorithm 1 Compute-Market-Prices

initialize $\mathbf{p}$ with $p_i \leftarrow \lceil w/s_i \rceil$.
repeat {for each i in turn}
  if p_i is *high* and $p_i - 1$ is not strictly *low* **then**
    decrement it ($p_i \leftarrow p_i - 1$)
  end if
until no decrement is effective

Claim 1. *A decrement of a price leaves all other high prices high.*

Proof. This is an immediate consequence of the WGS property. □

Lemma 2. *Algorithm 1 terminates after at most $n^2 w/s$ calls to the demand oracle.*

Proof. If a price p_i reaches 1, it is necessarily barely high. Further, by Claim 1 high prices remain high until decremented. If no price can be dropped, then this is the desired state, otherwise the price of some good is decremented. Since $p_i \geq 1$ for all goods i, there can be at most $(w/s)n$ decrements. Conceivably, we may have to repeatedly try all goods to find the one whose price reduces, and therefore the total number of queries is at most $n^2 w/s$. □

Lemma 3. *Decrementing the price of a good whose current price is strictly high and greater than $2ks_i$ increases its demand by at most 1.*

Proof. Let p_i denote the price of such a good after the decrement. Let Δ denote $x_i(p_i, \mathbf{p}_{-i}) - x_i(p_i + 1, \mathbf{p}_{-i})$. We claim that $\Delta \leq 1$. Note that $(1 + 1/p_i)^k \leq 1 + 1/s_i$ for $p_i \geq 2ks_i$. From k-smoothness, we have:

$$x_i(p_i, \mathbf{p}_{-i}) \leq x_i(p_i + 1, \mathbf{p}_{-i})\,(1 + 1/p_i)^k + 1.$$

Thus $x_i(p_i, \mathbf{p}_{-i}) \leq x_i(p_i + 1, \mathbf{p}_{-i})(1 + 1/s_i) + 1$. Therefore,

$$\Delta \quad \leq \quad x_i(p_i + 1, \mathbf{p}_{-i})/s_i + 1$$

But $x_i(p_i + 1, \mathbf{p}_{-i}) < s_i$ as $p_i + 1$ is strictly high. As Δ is an integer, $\Delta \leq 1$. □

Proof. (Lemma 4) If $p_i \leq k$ then the lemma follows trivially since the excess supply is always less than the total supply s_i and $s_i \leq ks_i/p_i + 1$ for $p_i \leq k$. Suppose that $p_i > k$. From k-smoothness, we have

$$x_i(p_i, \mathbf{p}_{-i}) \geq \frac{x_i(p_i - 1, \mathbf{p}_{-i}) - 1}{(1 + 1/(p_i - 1))^k} \geq \frac{x_i(p_i - 1, \mathbf{p}_{-i})}{(1 + 1/(p_i - 1))^k} - 1 = x_i(p_i - 1, \mathbf{p}_{-i})\,(1 - 1/p_i)^k - 1$$

Since p_i is barely high, $p_i - 1$ must be low, so $x_i(p_i - 1, \mathbf{p}_{-i}) \geq s_i$. Therefore $s_i - 1x_i(p_i, \mathbf{p}_{-i}) \leq s_i - s_i(1 - 1/p_i)^k + 1$. As $(1 - 1/p_i)^k \geq 1 - k/p_i$ for $p_i > k$, the excess supply is given by:

$$s_i - x_i(p_i, \mathbf{p}_{-i}) \leq s_i(1 - 1 + k/p_i) + 1 = ks_i/p_i + 1.$$

□

Lemma 4. *The excess supply of a good whose current price p_i is barely high and less than or equal to $2ks_i$ is at most $ks_i/p_i + 1$.*

Recall that the inefficiency of an allocation is defined as the ratio of the total (buyer and seller) discontent and the total (buyer and seller) value. For subsequent analysis, it will be helpful to rewrite the expression for Market Efficiency (Defn. 5) and specify the terms for buyers and sellers separately as:

$$1 - \frac{\sum_{b \in B} d_b((A_b, w_b - A_b \cdot \mathbf{p}), \mathbf{p}) + \sum_{g \in G} c_g p_g}{\sum_{b \in B} v_b(A_b^{opt}, m_b^{opt} \cdot \mathbf{p}) + \sum_{g \in G} s_g p_g}$$

where c_g denotes the number of unsold copies of g in the allocation, v_b denotes the upper value for buyer b and (A_b^{opt}, m_b^{opt}) is an optimal allocation (including money) for b with wealth w_b at prices $\mathbf{p}$.

Lemma 5. *The above allocation at the prices* $\mathbf{p}$ *computed by Algorithm 1 has efficiency* $1-O(k/r)$.

Proof. As each buyer receives its optimal allocation, there is no buyer discontent. The only discontent is the seller discontent. Let $\mathbf{p}$ denote the set of prices computed by the algorithm. We claim that a good i whose price p_i is greater than $2ks_i$ is correctly priced. For p_i is barely high, $p_i - 1$ is low, and by Lemma 3, the demand for such a good increases by at most 1 at each price decrement; therefore the demand exactly matches the supply at prices $\mathbf{p}$.

For a good i whose price is less than or equal to $2ks_i$, by Lemma 4, there are at most $ks_i/p_i + 1$ unsold copies. The total seller discontent is therefore at most

$$\sum_{i:p_i \le 2ks_i} (ks_i/p_i + 1)\, p_i \le \sum_{i:p_i \le 2ks_i} ks_i + \sum_{i:p_i \le 2ks_i} p_i \le k \sum_i s_i + 2k \sum_i s_i = 3k \sum_i s_i$$

Recall that the inefficiency of a market is defined as the ratio of the total (buyer and seller) discontent and the total (buyer and seller) value. The inefficiency of the allocation at prices $\mathbf{p}$ is then given by:

$$\frac{3k \sum_i s_i}{\sum_j w_j + \sum_i s_i p_i} \le \frac{3k \sum_i s_i}{\sum_j w_j} = O(k/r).$$

The efficiency is thus $1 - O(k/r)$. □

Let $\epsilon_i = 1/(8ks_i)$, $t = \sum_i s_i$ and $\epsilon = \min_i \epsilon_i$.

Theorem 6. *Algorithm 1 can be modified so that it makes at most* $O(nkt \log w)$ *demand oracle queries.*

Proof. Instead of considering all prices in the set $\{1, \ldots, \lceil w/s_i \rceil\}$, one can consider multiplicatively growing values of p_i from the set $\mathbb{I}^i_+ = \{1, 2, \ldots, \lceil 1/\epsilon_i \rceil\} \cup \{\lfloor (1+\epsilon_i)^j \rfloor : j > 0, (1+\epsilon_i)^j \le \lceil w/s_i \rceil\}$. Note is that if at price $\lfloor (1+\epsilon_i)^{j+1} \rfloor$, there is an excess supply of good i, then at price $\lfloor (1+\epsilon_i)^j \rfloor > 8ks_i$, the excess supply reduces by at most 1, due to the following argument. Let $p_i = \lfloor (1+\epsilon_i)^j \rfloor > 8ks_i$ and $p_i + h = \lfloor (1+\epsilon_i) \rfloor^{j+1}$ so that $h \le (1+\epsilon_i)^{j+1} - \big((1+\epsilon_i)^j - 1\big) = \epsilon_i (1+\epsilon_i)^j + 1 \le 2\epsilon_i (1+\epsilon_i)^j$. Similarly, $p_i = \lfloor (1+\epsilon_i)^j \rfloor \ge (1+\epsilon_i)^j - 1 \ge (1+\epsilon_i)^j/2$. Therefore, $h/p_i \le 4\epsilon_i$ and $(1 + h/p_i)^k \le (1+4\epsilon_i)^k \le 1 + 1/s_i$ for $4\epsilon_i \le 1/(2ks_i)$. From this point on, an argument identical to the one used in the proof of Lemma 3 works. It follows that for a good i whose price is greater than $8ks_i$ upon the termination of the algorithm, the demand equals the supply and there are no unsold copies.

Since $|\mathbb{I}^i_+| \le 8ks_i(1+\log w)$, the complexity becomes $8nk(\sum_i s_i)(1+\log w) = O(nkt \log w)$ oracle calls. □

When t is too large one can reduce the accuracy of the low test without significantly affecting the quality of the approximation.

Definition 7. *Price* p_i *is* near-low *if the excess demand for good* i *at* $(p_i, \mathbf{p}_{-i})$ *is at most* s_i/r.

The test in the algorithm (in line 3) is changed from a test for a low price to a test for a near-low price. Further, for all i, ϵ_i is set to $1/(8kr)$. This may create a further $O(s_i/r)$ excess supply of good i, but it is clear that this does not affect the approximation quality. The complexity becomes $O(krn^2 \log w)$ oracle calls.

k-smooth markets obeying the relaxed WGS property As mentioned earlier, many widely studied utility functions that obey the WGS property in the divisible case do not do so in the indivisible setting; this includes a range of CES utility functions. We illustrate this counter-intuitive behavior via the following example.

Example 1. *Consider a buyer with initial wealth* \$29 *in a market with two goods* g_1, g_2*, with demand* n_i *for good* g_i*,* $i = 1, 2$*. The buyer has a CES utility function of the form* $u(n_1, n_2) = (1 + \epsilon)\sqrt{n_1} + \sqrt{n_2}$*, for some small* $\epsilon > 0$*. Suppose* $p_1 = 10, p_2 = 10$*. At these prices the buyer's optimal bundle is* $(1, 1)$*. However, at* $p_1 = 10$*,* $p_2 = 9$ *the optimal bundle for this buyer is* $(2, 1)$*. This increase in demand for* g_1 *on having reduced the price of* g_2 *is a violation of the WGS property.*

We relax the WGS property, analogously to k-smoothness, as follows:

Definition 8. *A market satisfies the* k*-relaxed WGS property if for all* $\mathbf{p}'_{-i} \leq \mathbf{p}_{-i}$*,*

$$x_i(p_i, \mathbf{p}'_{-i}) \leq x_i(p_i, \mathbf{p}_{-i})(1 + k/p_i).$$

Theorem 7. *In a* k*-smooth market that satisfies the* k*-relaxed WGS property, a modification of Algorithm 1 computes prices* $\mathbf{p}$ *at which making the corresponding natural allocation has efficiency* $1 - O(k/r)$*.*

Proof. The idea is to modify Algorithm 1 to set aside some number of copies of each good to account for the relaxed WGS property. These set aside copies are used in case decreasing the prices of other goods causes the demand for good i to increase. Note that the maximum amount by which the demand for good i can increase as the prices of other goods decrease (the relaxed WGS property) is upper bounded by ks_i/p_i.

For $p_i > ks_i$, $ks_i/p_i < 1$ and no copies need to be set aside. For $p_i \leq ks_i$, we change line 3 of Algorithm 1 and decrement p_i only if the decrement leaves at least ks_i/p_i "spare" copies of good i (in surplus). The number of unsold copies of good i such that $1 < p_i \leq ks_i$ is upper bounded by $ks_i/p_i + ks_i/p_i + 1$ (by an argument similar to the one presented in the proof of Lemma 4, given in the appendix).

Thus, the efficiency of the algorithm is as before: $1 - O(k/r)$. □

Note that the algorithm needs to know the parameter k in markets that obey the k-relaxed WGS property.

4 Hardness

4.1 A Variant of SAT

We will use the following Max-3SAT-B problem as the underlying NP-complete language:

Input. A CNF formula with n variables in which every clause contains at most 3 literals, and every variable appears in at most a bounded number (B) of clauses.

Output. The maximum number of clauses can be satisfied simultaneously by some assignment to the variables.

Arora et al. [1] prove the following:

Theorem. *For some fixed $\epsilon > 0$, it is NP-hard to distinguish between satisfiable 3CNF-B formulas, and 3CNF-B formulas in which at most a $(1-\epsilon)$-fraction of the clauses can be satisfied simultaneously.*

Our reduction will be from the balanced 3SAT-3 problem.

Definition 9. (Balanced Max-3SAT-3 problem.) Input. *A CNF formula with n variables in which every clause contains at least 2 and at most 3 literals (a literal is a boolean variable in either positive or negative form), and every literal appears in exactly 3 clauses.*

Output. *The maximum number of clauses that can be satisfied simultaneously by some assignment to the variables.*

Theorem 8. The balanced Max-3SAT-3 problem is APX-Hard: *For some fixed $\delta > 0$, it is NP-hard to distinguish between satisfiable 3CNF-3 formulas, and 3CNF-3 formulas in which at most a $(1-\delta)$-fraction of the clauses can be satisfied simultaneously.*

Proof. Given an instance ψ of a Max-3SAT-B problem, we reduce it to an instance ϕ of the balanced Max-3SAT-3 problem in polynomial time.

Consider any variable x and let b be the number of occurrences of x in ψ. Let b_+ and b_- denote the number of positive and negative occurrences of x, respectively. The first step of the reduction is to add $|b_+ - b_-|$ *balance* clauses, where each balance clauses reduces the imbalance between the positive and negative literals by 1. Without loss of generality, suppose that $b_+ > b_-$; we add $b_+ - b_-$ clauses of the form $(\overline{x} \vee \overline{x} \vee x)$. The addition of these clauses causes the boolean formula to be balanced.

Suppose that there are b_x occurrences of x (and therefore of $\overline{x}$) in the new formula. Pair each occurrence of x with an occurrence of $\overline{x}$. Since the formula is balanced, all occurrences are paired. Replace the occurrence of x in the ith pair with a fresh variable x_i, for $0 \leq i \leq b_x - 1$, and add $2b_x$ clauses $(x_i \vee \overline{x}_{i+1}), (\overline{x}_i \vee x_{i+1})$, where $i+1$ is computed mod b_x. The new clauses are satisfied if and only if $x_i = x_{i+1}$ for every i. Now each literal appears exactly 3 times, and no literal appears more than once in the same clause. The above reduction has the following properties:

- The reduction takes polynomial time.
- If ψ is satisfiable, then so is ϕ.
- The number of clauses increases by at most a constant multiplicative factor. Let m, m_1, m_2 be the number of clauses in the formula ψ, after the first step, and in the final formula ϕ respectively. We have $m_2 \leq m_1 + 2\sum_x b_x$, yielding $m_2 = O(m_1)$ and $m_1 \leq 4m = O(m)$. Thus $m_2 = O(m)$ as desired.
- The number of unsatisfiable clauses decreases by at most a constant multiplicative factor. We need to argue that if ϵm clauses in ψ are not satisfied for every truth assignment to ψ, then δm_2 clauses in ϕ are not satisfied for every truth assignment to ϕ, for a suitable constant $\delta > 0$. We prove the contrapositive. Assume that fewer than δm_2 clauses in ϕ are not satisfied. Then we construct an assignment σ_ψ from an assignment σ_ϕ as follows: for a boolean variable x in ψ, if each one of $x_i, 0 \leq i \leq b_x - 1$ has the same truth value in σ_ϕ, then set x to that value. If however, some x_is are true while others are false, pick the majority of the two as the value of x. Each unsatisfied clause in ϕ yields at most B unsatisfied clauses in ψ.

The last two properties imply that if a ϵ-fraction of the clauses of ψ are not satisfiable, then an $\frac{\epsilon m}{Bm_2}$-fraction of the clauses of ϕ are not satisfiable; setting $\delta = \frac{\epsilon m}{Bm_2}$ suffices. □

4.2 The Reduction

The construction has two parts. The core is a market with $O(1)$ copies of a constant number of goods per variable in the 3-CNF formula, which achieves $\Theta(1)$-hardness: no discontent if the formula is satisfiable and $\Theta(\alpha)$ discontent if an α fraction of the clauses are not satisfied. The key building block is a device to choose prices; this device relies on the (slight) failure to satisfy the WGS property to enable this choice.

The remainder of the structure introduces more copies of each good so as to satisfy the k-WGS and k-smoothness properties. The basic approach is to introduce a separate wealthy buyer for each good. The main difficulty is that for goods for which there is a real choice of prices, the demand of the wealthy buyer must be the same at each of these prices, which limits such a buyer's wealth. Overcoming this constraint for small k, $k = 1 + \epsilon$, $\epsilon > 0$ an arbitrarily small constant, is quite delicate.

We show the following:

Theorem 9. *There is a fixed $\beta > 0$ such that in a market with average wealth $\Theta(r)$, it is NP-hard to compute prices at which a $(1 - \beta k/r)$ approximation to the efficiency is possible.*

Notation: In the boolean formula, we let n and m denote the number of variables and the number of clauses, respectively. Note that in a balanced 3SAT-3 formula, $m = \Theta(n)$. For ease of exposition, we assume that all clauses contain exactly 3 literals. The construction is readily modified when some of the clauses contain 2 literals instead of 3.

Given a balanced 3SAT-3 problem, we construct a k-smooth market with $\Theta(nr^2)$ wealth that satisfies the k-relaxed WGS property, with $\Theta(r)$ wealth per item, in which it is NP-hard to find prices that support an allocation with less than $\Theta(nkr)$ discontent.

For each variable x in the boolean formula, we introduce a separate collection of gadget submarkets. Together, they ensure that for each variable either there is a pricing that corresponds to a truth assignment and which has 0 discontent, or there is $\Theta(kr)$ discontent. For each clause C, we introduce another collection of gadget submarkets that enforce a pricing in which satisfied clauses yield 0 discontent and unsatisfied ones produce $\Theta(kr)$ discontent. We show that the construction satisfies k-smoothness and relaxed k-WGS conditions later.

For each variable x in the boolean formula, we create goods $G_x, G_{\overline{x}}$ (corresponding to the positive and negative literals of x) and two helper goods F_x and H_x. For each clause $C = a \vee b \vee c$, we introduce a "discontent" good D_C. When no ambiguity arises, we let $G, \overline{G}, F, H, D$ denote $G_x, G_{\overline{x}}, F_x, H_x, D_C$ respectively. For the remainder of this section, we let p_E denote the price of good E.

Definition 10. *We call prices $p_D = 4r/5, p_F = 3r/4, p_H = r + 2, p_G, p_{\overline{G}} \in \{r, r + 1\}$ and $p_G \neq p_{\overline{G}}$* select prices[5]*. For goods G and $\overline{G}$, prices within λk of the select prices for a suitable constant $\lambda > 0$ are said to be* broadly select*.*

Our interpretation is that $x = \textit{True}$ if $p_G \leq r$ and $p_{\overline{G}} > r$ and $x = \textit{False}$ if $p_{\overline{G}} \leq r$ and $p_G > r$. The reduction uses copies of the following gadget submarkets:

[5]Strictly, this should be $p_D = \lceil 4r/5 \rceil$, $p_F = \lceil 3r/4 \rceil$.

1. Force-Price-Up (E, r'): The purpose of this device is to help ensure that $p_E \geq r'$ or there will be $\Theta(r')$ discontent. We create one copy of good E and one buyer with money $2r' - 1$ that desires only E.

   Claim 10 below specifies the functioning of this gadget precisely.

2. Force-Price-Down (E, r'): The purpose of this device is to help ensure that $p_E \leq r'$ or there will be an unsold copy of E in this submarket. We create one copy of good E and one buyer with money r' that desires only E.

   Claim 11 below specifies the functioning of this gadget precisely.

3. Force-Price-Sum-Up $(G, \overline{G}, F, r)$: The purpose of this device is to help ensure that $p_G + p_{\overline{G}} > 2r$, or there will be $\Theta(r)$ discontent.

   We create 1 copy each of G and $\overline{G}$ and 2 copies of a floor good F for this submarket. We also create 2 *low* buyers with money $2r$ each.

   One low buyer mildly prefers G over $\overline{G}$, and the other one the other way round. In addition, a low buyer also has a desire for the floor good F. If $(n_G, n_{\overline{G}}, n_F)$ denotes the number of copies of $G, \overline{G}$ and F received by the low buyers, then the first buyer has preference order $(1,1,0) \succ (1,0,1) \succ (0,1,1) \succ (0,0,2)$ and the other one has preferences $(1,1,0) \succ (0,1,1) \succ (1,0,1) \succ (0,2,0)$. A CES function of the form $(2+\epsilon)\sqrt{n_G} + 2\sqrt{n_{\overline{G}}} + \sqrt{n_F}$ ($\epsilon > 0$ small) for example, captures $(1,1,0) \succ (1,0,1) \succ (0,1,1) \succ (0,0,2)$.

4. Force-Clause-True (G_a, G_b, G_c, D, r): The purpose of this gadget is to help obtain $\Theta(r)$ discontent if the clause $C = a \vee b \vee c$ is unsatisfied when the prices of goods $G_x, G_{\overline{x}}$, with $x \in \{a, b, c\}$, are set so as to correspond to a truth assignment. We create one copy each of goods G_a, G_b, G_c and a *clausal* buyer with money $3r+2$ and utility $(1+2\epsilon)\sqrt{n_a} + (1+\epsilon)\sqrt{n_b} + \sqrt{n_c} + (1/2)\sqrt{n_D}$, where ϵ is a small constant, say $1/r$.

   Note that even though a clausal buyer has a desire for D_C, no copies of D_C are introduced in this submarket. Claim 24 below specifies the functioning of this gadget precisely.

5. Big-Buyer (E, r', p'): The purpose of this gadget is to ensure that the demand for E obeys the k-smoothness and the k-relaxed WGS conditions. In addition, it is also used to ensure that the price of good E (in case E is G or $\overline{G}$) is broadly select, or there will be $\Theta(kr')$ discontent. We add r' copies of good E and a single buyer B' with $r'p'$ money desiring only good E.

6. Undo-Pair $(G, \overline{G}, H, r)$: We introduce 1 copy of good H and 1 copy each of G and $\overline{G}$ and two buyers with $2r + 2$ money and utilities $\sqrt{n_G} + \sqrt{n_H}$, $\sqrt{n_{\overline{G}}} + \sqrt{n_H}$, respectively.

   Claim 14 below specifies the functioning of this gadget precisely.

Claim 10. *In submarket Force-Price-Up (E, r'), if $p_E < r'$ there is an excess demand of at least $\lfloor r'/p_E \rfloor$ within the submarket.*

Proof. At $p_E < r'$, the buyer with $2r' - 1$ money desires $\lfloor (2r'-1)/p_E \rfloor$ copies of E. Since the supply of E in this submarket is 1, the excess demand is at least $\lfloor (2r'-1)/p_E \rfloor - 1 = \lfloor (2r' - 1 - p_E)/p_E \rfloor$. For $p_E < r'$, $2r' - 1 - p_E \geq r'$ and the excess demand for E within this submarket is at least $\lfloor r'/p_E \rfloor$. □

Claim 11. *In submarket* Force-Price-Down (E, r')*, if* $p_E > r'$ *there is an excess supply of 1 within the submarket.*

Proof. At $p_E > r'$, the buyer with r' money cannot afford any copy of E and there is 1 unsold copy within the submarket. □

The following two claims are immediate.

Claim 12. *In submarket* Force-Price-Up (E, r')*, if* $p_E = r'$ *the demand for* E *and the supply for* E *in the submarket are equal.*

Claim 13. *In submarket* Force-Price-Down (E, r')*, if* $p_E = r'$ *the demand for* E *and the supply for* E *in the submarket are equal.*

Claim 14. *In submarket* Undo-Pair $(G, \overline{G}, H, r)$*, if goods* $G, \overline{G}, H$ *are at their select prices (*$p_H = r + 2$ *and* $p_G = r, p_{\overline{G}} = r + 1$ *or* $p_G = r + 1, p_{\overline{G}} = r$*) then the demand for the more expensive good of* $G, \overline{G}$ *(*$\overline{G}$ *say) is 2 while the demand for the less expensive (*G *say) is 1. The demand for* H *is 1 at either set of prices.*

Proof. Suppose $p_H = r + 2, p_G = r$ and $p_{\overline{G}} = r + 1$. The buyer with utility $\sqrt{n_G} + \sqrt{n_H}$ desires 1 copy of H and 1 copy of G at these prices. However, the buyer with utility $\sqrt{n_{\overline{G}}} + \sqrt{n_H}$ desires no copy of H and 2 copies of $\overline{G}$. □

The Construction: Let c_1, c_2, c_3 and c_4 be suitable constants (to be specified later), independent of k. The construction is described as follows:

1. For each variable x, we introduce:
   (a) k instances of Force-Price-Up (G, r) and k instances of Force-Price-Down $(G, r + 1)$,
   (b) k instances of Force-Price-Sum-Up $(G, \overline{G}, F, r)$,
   (c) $3k$ instances of Force-Price-Up $(F, 3r/4)$ and $3k$ instances of Force-Price-Down $(F, 3r/4)$,
   (d) k instances of Force-Price-Up $(H, r + 2)$ and k instances of Force-Price-Down $(H, r + 2)$,
   (e) 1 instance of each of Big-Buyer $(G, c_1(r + 1), r)$, Big-Buyer $(\overline{G}, c_1(r + 1), r)$, Big-Buyer $(F, c_2 r, 3r/4)$, Big-Buyer $(H, c_3(r + 1), r + 2)$, and
   (f) c_1 instances of Undo-Pair $(G, \overline{G}, H, r)$.
2. For each clause $C = a \vee b \vee c$, we introduce:
   (a) k instances of Force-Clause-True (G_a, G_b, G_c, D, r),
   (b) k instances of Force-Price-Up $(D, 4r/5)$ and $4k$ instances of Force-Price-Down $(D, 4r/5)$, and
   (c) 1 instance of Big-Buyer $(D, c_4 r, 4r/5)$.

Outline: First, we show that for each good, if it is not at its select price, there is $\Theta(kr)$ discontent (actually for G and $\overline{G}$, at their broadly select prices). We then show that for each good pair $G, \overline{G}$, if they are at broadly select prices and the associated F and H goods are at select prices, then exactly one of G and $\overline{G}$ has a price of at most r or there is $\Theta(kr)$ discontent. Following this, we show that for each clause for which all goods are at select or broadly select prices, either at least one G (or $\overline{G}$) good has price $\leq r$ or there is $\Theta(kr)$ discontent. The former situation corresponds to a satisfied clause. It is then straightforward to relate discontent $\Theta(\beta kr)$ to $\Theta(\alpha n)$ unsatisfied clauses, for α, β fixed constants.

Claim 15. *If the price of good D is not $4r/5$, there is at least $\Theta(kr)$ discontent in the overall market, if $r > 10$.*

Proof. If $p_D > 4r/5$, by Claim 11, there is an excess supply of $4k$ within copies of the Force-Price-Down $(D, 4r/5)$ submarket. At this price, the clausal buyer in the relevant Force-Clause-True submarket desires at most 3 copies of D (since $4 \cdot 4r/5 > 3r + 2$ for $r > 10$). The demand from the k copies of the Force-Clause-True submarket is therefore at most $3k$ and there is an overall excess supply of at least $4k - 3k = k$, leading to seller discontent of at least $\Theta(kr)$.

If $p_D < 4r/5$, by Claim 10, there is an excess demand of at least $k \left\lfloor \frac{4r/5}{p_D} \right\rfloor$ within copies of the Force-Price-Up $(D, 4r/5)$ submarket; also the Big-Buyer $(D, c_4 r, 4r/5)$ submarket can only create further excess demand. Since there is no supply of D in any other submarket, there is buyer discontent of at least $k \left\lfloor \frac{4r/5}{p_D} \right\rfloor p_D \geq 4kr/10 = \Theta(kr)$ □

The following claim is immediate.

Claim 16. *If $p_D = 4r/5$, the demand for D in Force-Price-Up, Force-Price-Down and Big-Buyer submarkets is equal to its supply.*

Claim 17. *If the price of good F is not $3r/4$, there is at least $\Theta(kr)$ discontent in the overall market.*

Proof. If $p_F > 3r/4$, by Claim 11, there is an excess supply of $3k$ within copies of the Force-Price-Down $(F, 3r/4)$ submarket. At this price, a low buyer in the Force-Price-Sum-Up $(G, \overline{G}, F, r)$ submarket desires at most 2 copies of F and there is an excess demand of at most 2. Since there are k instances of the Force-Price-Sum-Up submarket, the overall excess supply is at least $3k - 2k = k$ leading to seller discontent of at least $\Theta(kr)$.

If $p_F < 3r/4$, by Claim 10, there is an excess demand of at least $3k \left\lfloor \frac{3r/4}{p_F} \right\rfloor$ within copies of the Force-Price-Up $(F, 3r/4)$ submarket; also Big-Buyer $(F, c_2 r, 3r/4)$ can only create further excess demand. There are at most $2k$ (unsold) copies of good F available from the Force-Price-Sum-Up $(G, \overline{G}, F, r)$. Thus, there is unmet demand of at least $k(3 \left\lfloor \frac{3r/4}{p_F} \right\rfloor - 2) \geq k \left\lfloor \frac{3r/4}{p_F} \right\rfloor$, leading to an overall discontent of $k \left\lfloor \frac{3r/4}{p_F} \right\rfloor p_F \geq 3kr/8 = \Theta(kr)$. □

Claim 18. *If the price of good H is not $r + 2$, there is at least $\Theta(kr)$ discontent in the overall market, if $r + 3 \geq 2c_3$ and $c_3 \geq c_1 + 1$.*

Proof. If $p_H > r + 2$, by Claim 11, there is an excess supply of k within copies of the Force-Price-Down $(H, r + 2)$ submarket. The excess supply in the Big-Buyer market is at least $c_3(r + 1) -$

$\left\lfloor \frac{c_3(r+1)(r+2)}{r+3} \right\rfloor \geq c_3 - \frac{2c_3}{r+3} \geq c_3 - 1$. At this price, a buyer in the Undo-Pair $(G, \overline{G}, H, r)$ submarket desires at most 1 copy of H, leading to excess demand of at most c_1 within copies of this submarket. Thus, there are at least $k + c_3 - 1 - c_1 \geq k$ unsold copies of H leading to seller discontent of at least $k(r+2) = \Theta(kr)$.

If $p_H < r+2$, by Claim 10, there is an excess demand of at least $k \left\lfloor \frac{r+2}{p_H} \right\rfloor$ within copies of the Force-Price-Up $(H, r+2)$ submarket. There is further excess demand of at least $\left\lfloor \frac{c_3(r+1)(r+2)}{r+1} \right\rfloor - c_3(r+1) = c_3$ from the Big-Buyer $(H, c_3(r+1), r+2)$ submarket. Since there are at most c_1 (unsold) copies of good H available from the Undo-Pair submarket, the excess demand within copies of the Force-Price-Up and the Big-Buyer submarkets is at least $k \left\lfloor \frac{r+2}{p_H} \right\rfloor + c_3 - c_1 \geq k \left\lfloor \frac{r+2}{p_H} \right\rfloor$ for $c_3 \geq 1 + c_1$, leading to buyer discontent of at least $k \left\lfloor \frac{r+2}{p_H} \right\rfloor p_H \geq k(r+2)/2 = \Theta(kr)$. □

Claim 19. *The supply of goods G and $\overline{G}$ in the overall market is $c_1(r+2) + 6k$.*

Proof. The supply of good G is $c_1(r+1)$ from the Big-Buyer submarket, c_1 from the Undo-Pair submarkets, k from the Force-Price-Up submarkets, k from the Force-Price-Down submarkets, k from the Force-Price-Sum-Up submarkets and $3k$ from the Force-Clause-True submarkets. Thus, the total supply is $c_1(r+2) + 6k$. □

Claim 20. *If $p_G \geq r + \lambda k$, then there is an excess supply of at least k for G, leading to seller discontent of $\Theta(kr)$, for $r \geq 2\lambda k, \lambda \geq 8/c_1$.*

Proof. By Claim 19, the supply of G in the overall market is $c_1(r+2) + 6k$. At $p_G \geq r + \lambda k$, the demand for G from the Big-Buyer submarket is at most:

$$\begin{aligned} \frac{c_1 r(r+1)}{r+\lambda k} &= c_1(r+1)\frac{1}{1+\lambda k/r} \\ &\leq c_1(r+1)(1 - \lambda k/r + \lambda^2 k^2/r^2) \\ &\leq c_1(r+1)(1 - \lambda k/2r) \qquad \text{(since } r \geq 2\lambda k\text{)} \\ &\leq c_1(r+1) - c_1 \lambda k/2 \end{aligned}$$

Buyers in the relevant Force-Clause-True submarkets may desire at most $6k$ copies of G and the demand from the Force-Price-Sum-Up submarket at this price is at most $2k$. Finally, there is a demand of at most k from copies of the Force-Price-Up submarket. Thus, the overall demand is at most $c_1(r+1) - c_1\lambda k/2 + 6k + 2k + k = c_1(r+1) + k(9 - \lambda c_1/2)$. The excess supply is therefore at least $c_1(r+2) + 6k - c_1(r+1) - k(9 - \lambda c_1/2) = c_1 + k(\lambda c_1/2 - 3)$. For $\lambda \geq 8/c_1$, the seller discontent is at least $\Theta(kr)$. □

Claim 21. *If $p_G \leq r - \lambda k$, then there is buyer discontent of at least $\Theta(kr)$ in the Big-Buyer submarket, for $\lambda \geq 7/c_1 + 1$ and $r \geq 2\lambda k$.*

Proof. By Claim 19, the supply of G in the overall market is $c_1(r+2) + 6k$. At $p_G = r - \lambda k$ (assume λk is an integer), the demand for G from the Big-Buyer submarket is $\left\lfloor \frac{c_1 r(r+1)}{r - \lambda k} \right\rfloor$. Since $\frac{c_1 r(r+1)}{r-\lambda k} = c_1(r+1)\frac{1}{1-\lambda k/r} \geq c_1(r+1) + c_1(r+1)\lambda k/r$, the demand from this submarket is at least $c_1(r+1) + c_1\lambda k$. Thus, the excess demand for G within the Big-Buyer submarket is at least:

$$\begin{aligned} c_1(r+1) + c_1\lambda k - (c_1(r+2) + 6k) &= c_1\lambda k - 6k - c_1 \\ &\geq c_1 k(\lambda - 1) - 6k \qquad \text{(since } k \geq 1\text{)} \\ &\geq k \qquad \text{(since } \lambda - 1 \geq 7/c_1\text{)} \end{aligned}$$

Since each unit of unmet demand leads to a discontent of $p_G = r - \lambda k$, the total buyer discontent is at least $k(r-\lambda k) = rk(1-\lambda k/r) = \Theta(kr)$. When the price halves, Big-Buyer's demand increases by a factor of 2 and therefore the buyer discontent at lower prices ($p_G \leq r - \lambda k$) is still $\Theta(kr)$. □

We say that a variable x is *mispriced* if either of the goods F_x and H_x is not at its select price (i.e. $p_{F_x} \neq 3r/4, p_{H_x} \neq r+2$) or if either of the goods $G_x, G_{\overline{x}}$ is not at its broadly select price. Note that by Claims 17, 18, 20 and 21, there is $\Theta(kr)$ discontent per mispriced variable x.

Discontent Sharing: We share the $\Theta(kr)$ discontent associated with a mispriced variable x uniformly among the Force-Clause-True submarkets in which x appears (either as a positive or a negative literal). Since a variable appears in exactly 6 clauses (in 3 as a positive literal and in 3 as a negative literal), there are a total of $6k$ Force-Clause-True submarkets among which the $\Theta(kr)$ discontent is shared. Suppose a is a literal of a mispriced variable x (i.e. $a = x$ or $a = \overline{x}$). The total discontent allocated to the k copies of Force-Clause-True (C, a, b, c) submarket is at least $k \cdot \Theta(\frac{kr}{6k}) = \Theta(kr)$.

Claim 22. *In submarket Force-Price-Sum-Up $(G, \overline{G}, F, r)$, if $p_G + p_{\overline{G}} \leq 2r$, $p_F = 3r/4$, and without loss of generality, $p_G \leq p_{\overline{G}}$, then the demand for G exceeds its supply in this submarket.*

Proof. Both low buyers desire at least one copy of G, creating a demand for at least 2 copies, while the supply is 1 copy. □

Claim 23. *If $p_F = 3r/4, p_H = r+2, p_D = 4r/5$ and goods G and $\overline{G}$ are at broadly select prices, then $p_G + p_{\overline{G}} > 2r$ or there is $\Theta(kr)$ discontent in the market.*

Proof. If $p_G + p_{\overline{G}} \leq 2r$, by Claim 22, buyers in the Force-Price-Sum-Up submarket will have an excess demand of k for each of G and $\overline{G}$. Without loss of generality, let $p_G \leq r$. If a Force-Clause-True submarket containing a copy of G contains another mispriced variable, then there is $\Theta(kr)$ shared discontent allocated to copies of this Force-Clause-True submarket. This discontent is further shared, equally among the literals that appear in the clause. Since there are at most 3 literals in any clause, each literal still receives $\Theta(kr)$ discontent. On the other hand, if each one of the literal goods appearing in this submarket is at its broadly select price, the buyer in this Force-Clause-True submarket desires at least 1 copy of G. Thus there is no excess supply of G within these submarkets. The buyer in an Undo-Pair having some desire for G desires at least 1 copy of G and the Big-Buyer seeks at least $c_1(r+1)$ copies of G. Thus, there is an overall excess demand of at least k for G in this case. Since any buyer that desires a copy of G in its optimal allocation and fails to receive it has $\Theta(r)$ discontent, there is an overall $\Theta(kr)$ buyer discontent in the market. □

Claim 24. *At broadly select prices, the buyer in Force-Clause-True (G_a, G_b, G_c, D, r) will buy one copy each of G_a, G_b, G_c, or there will be $\Theta(r)$ discontent, assuming $r \geq 10\lambda k$.*

Proof. If this buyer can afford these goods but does not receive them, it has $\Theta(r)$ discontent. Otherwise, its preferred basket includes the "discontent" good D, assuming $3r + 2 - 2(r+\lambda k) \geq \lceil 4r/5 \rceil \Rightarrow r/5 + 1 \geq 2\lambda k$; $r \geq 10\lambda k$ suffices. If it receives a copy of D, the Big-Buyer for D is short a copy, creating buyer discontent of $\Theta(r)$. If it does not receive a copy of D, it suffers buyer discontent $\Theta(r)$. □

Claim 25. *At select prices for all goods, the demand for goods F and H exactly equals their supply. Also, the total demand for goods $G, \overline{G}$ and D_C from all submarkets except the* Force-Clause-True *submarket exactly equals their supply in these submarkets.*

Proof. The demand for F at $p_F = 3r/4$ is of (a) $2k$ copies from the Force-Price-Sum-Up submarkets, (b) $3k$ copies from the Force-Price-Up submarkets, (c) $3k$ copies from the Force-Price-Down submarkets and (d) $c_2 r$ from the Big-Buyer submarket, which exactly equals the supply of F in the respective submarkets.

The demand for H at $p_H = r+2$ is of (a) k copies each from the Force-Price-Up and Force-Price-Down submarkets, (b) c_1 copies from the Undo-Pair submarket (Claim 14) and (c) $c_3(r+1)$ copies from the Big-Buyer submarket, which exactly equals the supply of H in the respective submarkets.

Other than the Force-Clause-True submarket, only Big-Buyer-D has a desire for D and when D is at its select price (i.e. $p_D = 4r/5$), the demand for D in this submarket is $c_4 r$, which exactly equals its supply.

When G and $\overline{G}$ are at their select prices (say $p_G = r$ and $p_{\overline{G}} = r+1$), the demand for G (in submarkets other than the Force-Clause-True submarket) is of (a) k copies each from the Force-Price-Up and Force-Price-Down submarkets, (b) k copies from the Force-Price-Sum-Up submarkets, (c) c_1 copies from the Undo-Pair submarkets (Claim 14) and (d) $c_1(r+1)$ from the Big-Buyer submarket. Thus, the total demand for G is $3k + c_1(r+2)$, which equals its total supply in these submarkets. The demand for $\overline{G}$ is same as that of G in these submarkets, except for the Undo-Pair submarkets and the Big-Buyer submarket. At $p_{\overline{G}} = r+1$, the demand for $\overline{G}$ is of $2c_1$ copies from the Undo-Pair submarkets (Claim 14) and of $c_1 r$ copies from the Big-Buyer submarket. The overall demand for $\overline{G}$ is still $3k + c_1(r+2)$, which equals its total supply in these submarkets. □

The constraints on r, λ, k are met by the following conditions $r \geq 10\lambda k, \lambda \geq 7/c_1 + 1, c_3 \geq c_1 + 1, r+3 \geq 2c_3, k \geq 1$.

Proof. (Theorem 9) Given a satisfying assignment σ, set prices as follows: $p_F = 3r/4, p_H = r+2$, if $\sigma(x) =$*True* then $p_{G_x} = r, p_{G_{\overline{x}}} = r+1$ (respectively, $\sigma(x) =$*False*, $p_{G_x} = r+1, p_{G_{\overline{x}}} = r$). For all clauses C, set $p_{D_C} = 4r/5$. At these prices, for each submarket, apart from the Force-Clause-True submarkets, by Claim 25, there is an allocation of goods introduced in that submarket to buyers in that submarket with 0 discontent. For a clause $C = a \vee b \vee c$, as at least one literal is true, in the corresponding Force-Clause-True (G_a, G_b, G_c, r) submarket, at least one of G_a, G_b, G_c is priced at r. Thus, the allocation to the buyer introduced in this submarket, of the goods introduced in this submarket, namely G_a, G_b, G_c has 0 discontent. Thus, all goods are sold and each buyer is optimally happy, leading to an overall discontent of 0.

Given a pricing, we construct a truth assignment for the formula as follows: if $p_{F_x} = 3r/4, p_{H_x} = r+2, p_{G_x} + p_{G_{\overline{x}}} > 2r$ and $p_{G_x} \leq r$, we set x to *True* (respectively if $p_{G_x} + p_{G_{\overline{x}}} > 2r$ and $p_{G_{\overline{x}}} \leq r$, we set x to *False*). Otherwise, we set the truth value arbitrarily. Suppose this leaves αm clauses unsatisfied. We show that there must be at least $\Theta(\alpha mkr)$ discontent. Consider one such unsatisfied clause $C = a \vee b \vee c$ in the formula under the given truth assignment. Let x, y and z represent the variables corresponding to literals a, b and c. There are three cases: (a) at least one of x, y and z is mispriced, (say there are $\alpha_1 m$ such clauses), or (b) x, y and z are not mispriced, but at least one of the literals a, b, c occurs with a literal d in another clause, and the variable corresponding to d is mispriced (say there are $\alpha_2 m$ such clauses), or (c) x, y and z are all correctly priced (i.e. not mispriced) and they each appear only in clauses with literals of other correctly priced variables (say there are $\alpha_3 m$ such clauses). For each of the $\alpha_1 m + \alpha_2 m$ clauses in cases (a) and (b), $\Theta(kr)$ discontent is allocated to the corresponding Force-Clause-True submarket (by sharing the discontent associated with at least $\max\{\alpha_1 m/6, \alpha_2 m/60\}$ mispriced variables). Finally, for the $\alpha_3 m$ clauses in case (c), there are two cases. If the price of the corresponding discontent good D is not $4r/5$, there is $\Theta(kr)$ discontent. Finally, if $p_D = 4r/5$, we argue that there must be an unsold copy of one of G_a, G_b, G_c, leading to a total of $\Theta(\alpha_3 mr)$ discontent as follows: with each of $p_{G_a}, p_{G_b}, p_{G_c}$ at at least $r+1$ (but no larger than $r + \lambda k$), the buyer in Force-Clause-True (C, G_a, G_b, G_c) with $3r+2$ money desires 1 copy each of G_a and G_b. Since each one of a, b, c appear only in clauses with correctly priced goods, no buyer in any other Force-Clause-True submarket can buy the one unsold copy of G_c from this submarket, leading to seller discontent of $k(r+1)$ from the k instances of the Force-Clause-True submarket.

Since the total wealth in the market is $\Theta(mr^2)$, the inefficiency is at least $\Theta(\alpha mkr)/\Theta(mr^2) = \Theta(\alpha k/r)$. Since we know it is hard to find a truth assignment with fewer than αm clauses unsatisfied, for some constant $\alpha > 0$, it must be hard to find a pricing with inefficiency less than $\beta k/r$ for a suitable constant $\beta > 0$. □

k-smoothness and relaxed k-WGS Let p_E denote the price of good E, $E \in \{F, H, D, G, \overline{G}\}$. At high enough prices, the only demand for E comes from the Big-Buyer that desires E. For ease of exposition (at the cost of a slightly clumsy notation), for the remainder of this section, we let Big-Buyer-E denote the Big-Buyer submarket corresponding to good E and we let x_E(all except Big-Buyer-E at $(p_E, \mathbf{p}_{-E})$) denote the demand for E from all submarkets other than Big-Buyer-E at prices $\mathbf{p} = (p_E, \mathbf{p}_{-E})$. When no ambiguity arises, we also let Big-Buyer-E denote the one buyer with money $r'p'$ introduced in Big-Buyer (E, r', p').

For goods F, H, D, G and $\overline{G}$, at $p_F > 2r, p_H > 2r+2, p_D > 3r+2, p_G > 3r+2$ and $p_{\overline{G}} > 3r+2$, there is no demand from any gadget submarket except the corresponding Big-Buyer submarket and the market obeys the k-smoothness and the relaxed k-WGS conditions, for $k \geq 1$, straightforwardly. At lower prices, a more careful analysis is required to demonstrate k-smoothness and relaxed k-WGS. The demand function for good E, $x_E(p_E, \mathbf{p}_{-E})$, obeys the k-smoothness condition if:

$$\text{demand for } E \text{ at } (p_E, \mathbf{p}_{-E}) \leq (\text{demand for } E \text{ at } (p_E+1, \mathbf{p}_{-E}))\,(1 + k/p_E).$$

For the above condition to hold, it suffices that:

$$\text{demand for } E \text{ at } (p_E, \mathbf{p}_{-E}) \leq (\text{demand for } E \text{ from Big-Buyer-}E \text{ at } (p_E+1, \mathbf{p}_{-E}))\,(1 + k/p_E).$$

The demand for E at price p_E from Big-Buyer-E with money w_B is $\lfloor w_B/p_E \rfloor$. Thus, to establish k-smoothness, it suffices to show that:

$$\begin{aligned}
\left\lfloor \frac{w_B}{p_E} \right\rfloor + x_E(\text{all except Big-Buyer-}E \text{ at } (p_E, \mathbf{p}_{-E})) &\leq \left\lfloor \frac{w_B}{p_E+1} \right\rfloor (1 + k/p_E) \\
(\Leftarrow)\ \frac{w_B}{p_E} + x_E(\text{all except Big-Buyer-}E \text{ at } (p_E, \mathbf{p}_{-E})) &\leq \left(\frac{w_B}{p_E+1} - 1 \right) (1 + k/p_E) \\
(\Leftrightarrow)\ 1 + \frac{w_B}{p_E(p_E+1)} + x_E(\text{all except Big-Buyer-}E \text{ at } (p_E, \mathbf{p}_{-E})) &\leq \frac{k}{p_E} \left(\frac{w_B}{p_E+1} - 1 \right) \\
(\Leftrightarrow)\ 1 + x_E(\text{all except Big-Buyer-}E \text{ at } (p_E, \mathbf{p}_{-E})) + \frac{k}{p_E} &\leq \frac{(k-1)w_B}{p_E(p_E+1)}
\end{aligned}$$

Rearranging the above inequality, we obtain a sufficient condition for k-smoothness, as expressed in the following lemma.

Lemma 26. *The demand for E at prices $(p_E, \mathbf{p}_{-E})$ in a market with* Big-Buyer-E *who has wealth w_B obeys the k-smoothness condition if:*

$$p_E + p_E\left(x_E(\textit{all except } \textsf{Big-Buyer-}E \textit{ at } (p_E, \mathbf{p}_{-E}))\right) + k \leq \frac{(k-1)w_B}{p_E+1} \tag{2}$$

Recall that we introduced one instance of each of Big-Buyer $(F, c_2 r, 3r/4)$, Big-Buyer $(H, c_3(r+1), r+2)$ and Big-Buyer $(D, c_4 r, 4r/5)$ in the construction. Also, note that we often use the inequality $2(k-1) \geq k$ for $k \geq 2$.

Lemma 27. *The demand for F, D and H obeys k-smoothness for $k \geq 2$, $r \geq 6$ and suitably large constants c_1, c_2, c_3 and c_4 with $c_3 \geq 24 + 8c_1$.*

Proof. First, let us consider good F. At $p_F > 2r$, no buyer other than Big-Buyer-F desires a copy of F, and the demand obeys k-smoothness for $k \geq 1$. At $p_F \leq 2r$, the demand for F may come from Force-Price-Sum-Up, Force-Price-Up and Force-Price-Down submarkets. The total wealth in these submarkets is $k(4r) + 3k(6r/4 - 1) + 3k(3r/4) \leq \alpha kr$, for a fixed constant α ($\alpha = 11$ works). Thus, $x_F(\text{all except Big-Buyer-}F \text{ at } (p_F, \mathbf{p}_{-F})) \leq \frac{\alpha kr}{p_F}$; on applying Lemma 26, we get the following sufficient condition for k-smoothness:

$$\begin{array}{lrcll}
 & p_F + \alpha kr + k & \leq & \frac{3c_2(k-1)r^2}{4(p_F+1)} & \\
(\Leftarrow) & 2r + \alpha kr + k & \leq & \frac{3c_2(k-1)r^2}{4(2r+1)} & (\text{since } p_F \leq 2r) \\
(\Leftarrow) & 2(k-1)r + 2\alpha(k-1)r + 2(k-1)r & \leq & \frac{c_2}{4}(k-1)r & (\text{since } k \geq 2, r \geq 1) \\
(\Leftrightarrow) & 4(4+2\alpha) & \leq & c_2 &
\end{array}$$

Similarly, for H, at $p_H > 2r+2$, only Big-Buyer-H desires H and the market obeys k-smoothness for $k \geq 1$. At $p_H \leq 2r+2$, there may be demand for H from Force-Price-Up, Force-Price-Down and Undo-Pair submarkets. The total wealth in these submarkets is $k(2r+3) + k(r+2) + c_1(4r+4) \leq 3k(r+2) + 4c_1(r+1) \leq \beta k(r+1) + 4c_1(r+1)$ for a fixed constant β ($\beta = 4$ works, assuming $r \geq 2$). Thus, $x_H(\text{all except Big-Buyer-}H \text{ at } (p_H, \mathbf{p}_{-H})) \leq \frac{\beta k(r+1) + 4c_1(r+1)}{p_H}$. Applying Lemma 26, we get the following sufficient condition for k-smoothness:

$$\begin{array}{lrcll}
 & p_H + \beta k(r+1) + 4c_1(r+1) + k & \leq & \frac{c_3(k-1)(r+1)(r+2)}{p_H+1} & \\
(\Leftarrow) & 2r + 2 + \beta k(r+1) + 4c_1(r+1) + k & \leq & \frac{c_3(k-1)(r+1)(r+2)}{2r+3} & (\text{since } p_H \leq 2r+2) \\
(\Leftarrow) & (k-1)(r+1)(2+2\beta+2) + 4c_1(r+1) & \leq & \frac{c_3}{2}(k-1)(r+1) & (\text{since } k \geq 2) \\
(\Leftrightarrow) & 2(4 + 2\beta + \frac{4c_1}{k-1}) & \leq & c_3 &
\end{array}$$

Finally, for D, at $p_D > 3r+2$, only Big-Buyer-D desires D and the market obeys k-smoothness for $k \geq 1$. At $p_D \leq 3r+2$, buyers in the corresponding Force-Clause-True, Force-Price-Up and Force-Price-Down submarket may also desire copies of D. The total wealth in these submarkets is $k(3r+2) + k(2\lceil 4r/5 \rceil - 1) + 4k(\lceil 4r/5 \rceil) \leq \gamma kr$ for a fixed constant γ ($\gamma = 9$ works, assuming $r \geq 6$). Applying Lemma 26, at $p_D \leq 3r+2$, we get the following sufficient condition for k-smoothness:

$$\begin{array}{lrcll}
 & p_D + \gamma kr + k & \leq & \frac{4c_4(k-1)r^2}{5(p_D+1)} & \\
(\Leftarrow) & 3r + 2 + \gamma kr + k & \leq & \frac{4c_4(k-1)r^2}{5(3r+3)} & (\text{since } p_D \leq 3r+2) \\
(\Leftarrow) & 4r + \gamma kr + kr & \leq & \frac{4c_4(k-1)r^2}{15(r+1)} & (\text{since } r \geq 2) \\
(\Leftarrow) & (k-1)r(4 + 2\gamma + 2) & \leq & \frac{2c_4}{15}(k-1)r & (\text{since } k \geq 2) \\
(\Leftrightarrow) & \frac{15}{2}(6 + 2\gamma) & \leq & c_4 &
\end{array}$$

□

Lemma 28. *The demand for G (and $\overline{G}$) obeys the k-smoothness condition, for $k \geq 6$, $r \geq 8$ and a suitably large constant c_1.*

Proof. We prove the statement for G. An identical argument works for $\overline{G}$. At $p_G > 3r+2$, the only demand for G comes from the Big-Buyer-G submarket and the market obeys k-smoothness for $k \geq 1$.

At $2r+2 < p_G \leq 3r+2$, the only buyers that may desire copies of G, other than the one in Big-Buyer-G, are the ones in the Force-Clause-True submarket, each of which can afford exactly one copy of G. Since there are at most $3k$ submarkets that contain a buyer that has a utility for G, $x_G(\text{all except Big-Buyer-}G \text{ at } (p_G, \mathbf{p}_{-G})) \leq 3k$. Applying Lemma 26, we get $p_G + p_G(3k) + k \leq \frac{(k-1)c_1 r(r+1)}{p_G+1}$ as a sufficient condition for k-smoothness. Since $p_G \leq 3r+2$, we get:

$$\begin{array}{rrcll}
 & (3r+2)(3k+1)+k & \leq & \frac{(k-1)c_1 r(r+1)}{3r+2} & \\
(\Leftarrow) & 3(r+1)7(k-1)+2(k-1) & \leq & \frac{(k-1)c_1 r(r+1)}{3r+2} & (\text{since } 2(k-1) \geq k \ \& \ k-1 \geq 1) \\
(\Leftarrow) & 23(k-1)(r+1) & \leq & \frac{8(k-1)c_1(r+1)}{26} & (\text{since } r \geq 8) \\
(\Leftrightarrow) & \frac{299}{4} & \leq & c_1 &
\end{array}$$

At prices $p_G \leq 2r+2$, the potential wealth that may be spent on G (other than the Big-Buyer) is: (a) $k(2r-1)$ from the Force-Price-Up submarket, (b) $k(r+1)$ from the Force-Price-Down submarket, (c) $4kr$ from the Force-Price-Sum-Up submarket, (d) $c_1(2r+2)$ from the Undo-Pair submarket and (e) $3k(3r+2)$ from the Force-Clause-True submarket. The total wealth that may be spent on good G (other than from the Big-Buyer-G submarket) is an upper bound on $p_G \cdot x_G(\text{all except Big-Buyer-}G \text{ at } (p_G, \mathbf{p}_{-G}))$; thus:

$$\begin{array}{rcl}
p_G \cdot x_G(\text{all except Big-Buyer-}G \text{ at } (p_G, \mathbf{p}_{-G})) & \leq & k(2r-1)+k(r+1)+4kr+3k(3r+2)+2c_1(r+1) \\
 & \leq & 16kr+6k+2c_1(r+1)
\end{array}$$

Applying Lemma 26 and using $p_G \leq 2r+2, k \geq 2$, it suffices that:

$$\begin{array}{rrcll}
 & (2r+2)+(16kr+2c_1(r+1)+6k)+k & \leq & \frac{(k-1)c_1 r(r+1)}{2r+3} & \\
(\Leftarrow) & 2(r+1)+16k(r+1)+2c_1(r+1) & \leq & \frac{(k-1)c_1 r(r+1)}{2r+3} & \\
(\Leftrightarrow) & 2+16k & \leq & c_1\frac{(k-1)r-(4r+6)}{2r+3} & \\
(\Leftarrow) & 18k & \leq & c_1\frac{(k-5-6/r)r}{2r+3} & (\text{since } k \geq 1) \\
(\Leftarrow) & 18k & \leq & c_1\frac{(k-\frac{23}{4})r}{2r+3} & (\text{since } r \geq 8) \\
(\Leftarrow) & 18k & \leq & c_1\frac{k}{24}\frac{r}{2r+3} & (\text{since } k \geq 6) \\
(\Leftarrow) & 18 & \leq & c_1\frac{1}{24}\frac{8}{19} & (\text{since } r \geq 8)
\end{array}$$

which is satisfied for a suitably large constant c_1. □

For the k-relaxed WGS condition, the permissible swing in demand changes from the bound of $\frac{k}{p_E}\left\lfloor\frac{w_B}{p_E+1}\right\rfloor$ (used in the k-smoothness analysis) to $\frac{k}{p_E}\left\lfloor\frac{w_B}{p_E}\right\rfloor$, which is larger than $\frac{k}{p_E}\left\lfloor\frac{w_B}{p_E+1}\right\rfloor$. Thus the previous analysis shows that the demand for $F, H, G, \overline{G}$ also obeys the k-relaxed WGS condition for $k \geq 6$.

Theorem 29. *For $r \geq 11k$, $k \geq 6$, there is a fixed $\beta > 0$ such that in a market with average wealth $\Theta(r)$ that obeys the k-smoothness and the relaxed k-WGS properties, it is NP-hard to compute prices at which a $(1-\beta k/r)$ approximation to the efficiency is possible.*

Comment 3. *The construction still works if buyers can be allocated arbitrary allocations. (This entails verifying that the claims of $\Theta(kr)$ discontent continue to apply).*

Comment 4. *A desire for money can also be introduced. The one difficulty occurs with the* Force-Price-Up *submarket; to create sufficient discontent on a low price, we introduce a discontent good, whose select price is $r/2$. This price is forced to be within $\pm\Theta(k)$ of $r/2$ by the* Big-Buyer *for this good (or there is $\Theta(kr)$ discontent).*

We would like to extend this result to arbitrarily small $k > 1$, which we can do. As the solution is somewhat painstaking, we only sketch it. The idea is to introduce a small collection of new buyers of G (and $\overline{G}$) with wealth $\Theta(c_1 r^2 \log c_1)$ such that Big-Buyer-G together with the new buyers assure the k-smoothness of new buyers, Big-Buyer-G ensures the k-smoothness of all submarkets except the Undo-Pair submarket and the new buyers ensure the k-smoothness of the c_1 buyers that desire G introduced in the Undo-Pair submarket. We show the result for $k = 1 + 4\alpha$, for an arbitrarily small but fixed $\alpha > 0$. Now, when defining the construction, rather than take $c_i k = c_i(1 + 4\alpha)$ copies of a submarket, for $1 \le i \le 4$, we use $c_i \lceil k \rceil = 2c_i$ copies. It is not hard to extend Lemma 27 to $k = 1 + 4\alpha$, where the constants c_1, c_2, c_4 are now $\Theta(1/\alpha)$ and c_3 is $\Theta(1/\alpha^2)$. The extension of Lemma 28 is more difficult.

Let $a = 2^h$, where h is a parameter to be specified later. We introduce buyers with money $4r, 4r + 1, 4r + 2, \ldots, ar - 1$. All these buyers desire only good G and have no desire for money. They are conceptually partitioned into $h - 2$ *levels*, 0 through $h - 3$. Among the newly introduced buyers, a buyer with wealth w such that $\frac{ar}{2^{j+1}} \le w < \frac{ar}{2^j}$ is said to be in level j. For each amount of wealth w in level j, there are 4^j identical buyers with that amount of wealth. Note that the average money with the buyers in a level and the level's length halve in successive levels, but since there are 4 times as many buyers with each amount of wealth, the total wealth in each level is (up to rounding differences) the same. There is one caveat: if $\left\lfloor \frac{w}{r} \right\rfloor \ne \left\lfloor \frac{w}{r+1} \right\rfloor$, no buyers of wealth w are introduced. This excludes buyers for wealth amounts $gr, gr + 1, \ldots, gr + (g - 1)$, for $4 \le g < a$. Note that this ensures that each newly introduced buyer has the same demand for G at prices r and $r + 1$.

Since we will be focusing only on good G for the remainder of this section, we drop the subscript G wherever possible and convenient. Let $x^U(p_G, \mathbf{p}_{-G})$ denote the demand of the Undo buyers for G at prices $(p_G, \mathbf{p}_{-G})$. Let $x^N(p_G)$ denote the demand of the new buyers for G at price p_G (note that the prices $\mathbf{p}_{-G}$ do not affect this demand). Let $x^B(p_G)$ denote the demand of Big-Buyer-G for G at price p_G. Finally, let $\widehat{x}(p_G, \mathbf{p}_{-G})$ denote the demand of the remaining buyers for G at prices $(p_G, \mathbf{p}_{-G})$. The k-smoothness bound will follow from the following three bounds, which we show:

$$x^B(p_G) - x^B(p_G + 1) \le \frac{1 + \alpha}{p_G} x^B(p_G + 1) \tag{3}$$

$$x^N(p_G) - x^N(p_G + 1) \le \frac{3\alpha}{p_G} x^B(p_G + 1) + \frac{1}{p_G} x^N(p_G + 1) \tag{4}$$

$$\widehat{x}(p_G, \mathbf{p}_{-G}) - \widehat{x}(p_G + 1, \mathbf{p}_{-G}) + x^U(p_G, \mathbf{p}_{-G}) - x^U(p_G + 1, \mathbf{p}_{-G}) \le \frac{3\alpha}{p_G} x^N(p_G + 1) \tag{5}$$

Claim 30. *Equations 3, 4 and 5 together imply $(1 + 4\alpha)$-smoothness.*

Proof. Summing Equations 3, 4 and 5, we obtain the following inequality:

$$\begin{aligned} x(p_G, \mathbf{p}_{-G}) - x(p_G+1, \mathbf{p}_{-G}) &\le \frac{1+4\alpha}{p_G} x^B(p_G+1) + \frac{1+3\alpha}{p_G} x^N(p_G+1) \\ &\le \frac{1+4\alpha}{p_G} x(p_G+1, \mathbf{p}_{-G}) \end{aligned}$$

Upon rearranging the above inequality, we get $x(p_G, \mathbf{p}_{-G}) \le x(p_G+1, \mathbf{p}_{-G})(1+\frac{1+4\alpha}{p_G})$, which implies $(1+4\alpha)$-smoothness. □

Claim 31. *Equation 3 holds if $p_G < ar$, where $a = \sqrt{\alpha c_1}$ and c_1 is an integer multiple of a.*

Proof. Let w_B denote the wealth of the Big-Buyer. By construction, $w_B = c_1 r(r+1)$. For positive integers a and q, $a \bmod q = a - q\lfloor a/q \rfloor < q$. Using this observation, we can verify Equation 3 since:

$$\begin{aligned} & w_B - (p_G+1)\left\lfloor \frac{w_B}{p_G+1} \right\rfloor &\le\ & p_G \\ (\Rightarrow)\quad & \frac{w_B}{p_G} - \frac{p_G+1}{p_G}\left\lfloor \frac{w_B}{p_G+1} \right\rfloor &\le\ & 1 \\ (\Rightarrow)\quad & \left\lfloor \frac{w_B}{p_G} \right\rfloor - \left\lfloor \frac{w_B}{p_G+1} \right\rfloor &\le\ & 1 + \frac{1}{p_G}\left\lfloor \frac{w_B}{p_G+1} \right\rfloor \\ (\Leftrightarrow)\quad & x^B(p_G) - x^B(p_G+1) &\le\ & 1 + \frac{1}{p_G} x^B(p_G+1). \end{aligned}$$

Thus, it suffices to show that $\frac{\alpha}{p_G} x^B(p_G+1) \ge 1$ in order to prove the claim. Indeed, we have:

$$\begin{aligned} \frac{\alpha}{p_G} x^B(p_G+1) &= \frac{\alpha}{p_G}\left\lfloor \frac{c_1 r(r+1)}{p_G+1} \right\rfloor \\ &\ge \frac{\alpha}{ar-1}\left\lfloor \frac{c_1 r(r+1)}{ar} \right\rfloor && \text{(since } p_G < ar\text{)} \\ &= \frac{\alpha c_1 (r+1)}{a(ar-1)} && \text{(since } c_1 \bmod a = 0\text{)} \\ &= \frac{a(r+1)}{ar-1} && \text{(since } a = \sqrt{\alpha c_1}\text{)} \\ &\ge 1. \end{aligned}$$

□

Note that the demand from the remaining buyers, $\widehat{x}(p_G, \mathbf{p}_{-G}) \le 2 \cdot \frac{16r+6}{p_G} \le \frac{34r}{p_G}$ for $r \ge 6$. Also note that $x^U(p_G, \mathbf{p}_{-G}) \le \frac{c_1(2r+2)}{p_G} \le \frac{3c_1 r}{p_G}$ for $r \ge 2$. Thus, to show Equation 5, assuming $r \ge 6$, it suffices to show:

$$\frac{34r}{p_G} + \frac{3c_1 r}{p_G} \le \frac{3\alpha}{p_G} x^N(p_G+1)$$

$$\text{or} \qquad (34 + 3c_1) r \le 3\alpha x^N(p_G+1) \tag{6}$$

It remains to compute upper and lower bounds on $x^N(p_G)$, and then it will be straightforward to check Equations 4 and 6.

Claim 32. *Let* $\left\lfloor \frac{ar-1}{i+1} \right\rfloor < p_G \le \left\lfloor \frac{ar-1}{i} \right\rfloor$*. Then*

$$
\begin{array}{rcllcl}
x^N(p_G) & \ge & \frac{ar}{2} + 4\cdot\frac{ar}{4} + \ldots + 2^{\lfloor \log i \rfloor -1}\frac{ar}{2} & + & 4^{\lfloor \log i \rfloor}\left(\frac{ar}{2^{\lfloor \log i \rfloor}} - 1 - (p_G - 1)\right) \\
& + & \frac{ar}{2} + 4\cdot\frac{ar}{4} + \ldots + 2^{\lfloor \log i \rfloor -2}\frac{ar}{2} & + & 4^{\lfloor \log i \rfloor -1}\left(\frac{ar}{2^{\lfloor \log i \rfloor -1}} - 1 - (2p_G - 1)\right) \\
& + & \ldots & & \\
& + & \frac{ar}{2} + \ldots + \ldots + 2^{\lfloor \log \frac{ar}{gp_G} \rfloor -2}\frac{ar}{2} & + & 4^{\lfloor \log \frac{ar}{gp_G} \rfloor -1}\left(\frac{ar}{2^{\lfloor \log \frac{ar}{gp_G} \rfloor -1}} - 1 - (gp_G - 1)\right) \\
& + & \ldots & & \\
& + & (ar-1) - (ip_G - 1) & & \\
& - & i[a + (a-1) + \ldots & + & a/2 + 1] \\
& - & 4\lfloor i/2 \rfloor[a/2 + (a/2-1) + \ldots & + & a/4 + 1] \\
& - & \ldots & & \\
& - & 4^{\lfloor \log i \rfloor}[\left\lfloor \frac{a}{2^{\lfloor \log i \rfloor}} \right\rfloor + \ldots + \ldots & + & \left\lfloor \frac{a}{2^{\lfloor \log i \rfloor + 1}} \right\rfloor + 1]
\end{array}
$$

Proof. In turn, we count the number of buyers seeking at least one copy of G, two copies, three copies, etc., ignoring the absence of buyers with wealth w such that $\lfloor \frac{w}{r} \rfloor \ne \left\lfloor \frac{w}{r+1} \right\rfloor$ (i.e. assuming there are buyers associated with such w). We then subtract an overestimate of the demand from these excluded buyers. □

Claim 33. *Equation 6 holds if* $r \ge 4a$ *and* $\log c_1 \ge 8 + \frac{128}{3\alpha^2} + \log\frac{1}{\alpha}$*, where* $a = \sqrt{\alpha c_1}$*.*

Proof. By Claim 32, $x^N(p_G + 1)$ can be expressed with the same formula but with i' replacing i, where $i' = i$ if $p_G < \lfloor \frac{ar-1}{i} \rfloor$ and $i' = i - 1$ if $p_G = \lfloor \frac{ar-1}{i} \rfloor$.

We verify Equation 6 for $p_G + 1 = 4r - 1$; it then follows for all smaller p_G as $x^N(p_G + 1)$ only grows as p_G shrinks. At $p_G + 1 = 4r - 1$, $i' = a/4$.

$$
\begin{aligned}
x^N(4r-1) \ \ge\ & \frac{ar}{2}(a/4-1) + \frac{ar}{2}(a/8-1) + 2\frac{ar}{2}(a/16-1) + \ldots + \frac{a}{8}\frac{ar}{2}(2-1) \\
& - \left[\frac{a^2}{16}\left(\frac{3a}{2}+1\right) + \frac{a^2}{16}\left(\frac{3a}{4}+1\right) + \ldots + \frac{a^2}{16}\left(\frac{3a}{a/2}+1\right)\right] \\
\ge\ & \frac{a^2 r}{8} + \frac{a^2 r}{16}(\log a - 3) - \frac{a^2 r}{8} \\
& - \left[\frac{a^2}{16}\left(\frac{3a}{2}+1\right) + \frac{a^2}{16}\left(\frac{3a}{4}+1\right) + \ldots + \frac{a^2}{16}\left(\frac{3a}{a/2}+1\right)\right] \\
\ge\ & \frac{a^2 r}{16}(\log a - 3) - \frac{3a^3}{16} - \frac{a^2}{16}(\log a - 1)
\end{aligned}
$$

We need $(34 + 3c_1)r \le 3\alpha x^N(4r-1)$. With $a = \sqrt{\alpha c_1}$ and $c_1 \ge 34$, it suffices that:

$$
\begin{array}{rrcll}
 & 4a^2 & \le & \frac{3\alpha^2}{16}\left(a^2(\log a - 3) - \frac{a^2}{r}(3a + \log a)\right) & \\
\text{or,} & 4a^2 & \le & \frac{3\alpha^2}{16}\left(a^2(\log a - 3) - \frac{4a^3}{r}\right) & \text{(if } a \ge \log a\text{)} \\
\text{or,} & 4 & \le & \frac{3\alpha^2}{16}\left(\log a - 3 - \frac{4a}{r}\right) &
\end{array}
$$

Thus it suffices that:

$$\begin{array}{rrcll} & \frac{64}{3\alpha^2}+3+\frac{4a}{r} & \leq & \log a & \\ \text{or,} & \frac{64}{3\alpha^2}+4 & \leq & \log a & (\text{if } r\geq 4a) \\ \text{or,} & 8+\frac{128}{3\alpha^2}+\log\frac{1}{\alpha} & \leq & \log c_1 & (\text{since } a=\sqrt{\alpha c_1}) \end{array}$$

□

Claim 34. *Equation 4 holds if* $r \geq (3a+10)/4$.

Proof. Temporarily, we pretend there are no excluded buyers w. Let $\left\lfloor \frac{ar-1}{i+1} \right\rfloor < p_G \leq \left\lfloor \frac{ar-1}{i} \right\rfloor$. We consider the buyers with demand f at price p_G+1, for $f \neq \left\lfloor \frac{i}{2^d} \right\rfloor$, $d \geq 0$. Clearly, for $a/2^{g+1} \leq f < a/2^g$, there are $4^g(p_G+1)$ such buyers. Hence their demand, divided by p_G+1 is exactly $4^g f$. Now, observe that there are $4^g f$ buyers with demand f at price p_G+1 and demand $f+1$ at price p_G. But this is the increase in demand at price p_G compared to price p_G+1 contributed by all buyers with demand f at price p_G+1.

The unaccounted increased demand at price p_G is at most $i+4\lfloor i/2\rfloor+\ldots+4^{\lfloor \log i\rfloor}\lfloor i/2^{\lfloor \log i\rfloor}\rfloor \leq 2i^2$.

There is also demand due to the excluded buyers with wealth w such that $\left\lfloor \frac{w}{r} \right\rfloor \neq \left\lfloor \frac{w}{r+1} \right\rfloor$ which is not really present at p_G+1; its (lack of) contribution has to be accounted for. But it amounts to at most:

$$\begin{array}{rl} & a+(a-1)+\ldots+(a/2+1) \\ & +\,4\,(a/2+(a/2-1)+\ldots+(a/4+1)) \\ & +\ \ldots \\ & +\,4^{\lfloor \log i\rfloor}\left(a/2^{\lfloor \log i\rfloor}+\ldots+a/2^{\lfloor \log i\rfloor+1}+1\right) \\ \leq & \frac{a}{4}\left(\frac{3a}{2}+1\right)+\frac{4a}{8}\left(\frac{3a}{4}+1\right)+\ldots+\frac{4^{\lfloor \log i\rfloor}a}{2^{\lfloor \log i\rfloor+2}}\left(\frac{3a}{2^{\lfloor \log i\rfloor+1}}+1\right) \\ \leq & \frac{3a^2}{8}(\lfloor \log i\rfloor+1)+2^{\lfloor \log i\rfloor-1}a \end{array}$$

The demand due to Big-Buyer-G at p_G+1 is $\left\lfloor \frac{c_1 r(r+1)}{p_G+1} \right\rfloor$. We need

$$2i^2+\frac{\frac{3}{8}a^2\left(\lfloor \log i\rfloor+1\right)+2^{\lfloor \log i\rfloor-1}a}{p_G+1} \leq 3\alpha\left\lfloor \frac{c_1 r(r+1)}{p_G+1} \right\rfloor/p_G$$

We note that $\left\lfloor \frac{c_1 r(r+1)}{p_G+1} \right\rfloor \geq \frac{c_1 r^2}{p_G+1}$ as $p_G+1 \leq c_1 r$. Thus it suffices to show that:

$$2i^2(p_G+1)p_G+\left(\frac{3a^2}{8}(\log i+1)+\frac{ai}{2}\right)p_G \leq 3\alpha c_1 r^2.$$

We recall that $\alpha c_1 = a^2$ and $p_G \leq \frac{ar-1}{i}$. So it suffices that:

$$\begin{array}{rrcl} & 2(ar-1+i)(ar-1)+\left(\frac{3a^2}{8}(\log i+1)+\frac{ai}{2}\right)\left(\frac{ar-1}{i}\right) & \leq & 3a^2r^2 \\ \text{or,} & 2a^2+\frac{2ia}{r}+\frac{3a^3}{8ir}\log i+\frac{3a^3}{8ir}+\frac{a^2}{2r} & \leq & 3a^2 \\ \text{or,} & \frac{2i}{ar}+\frac{3a\log i}{8ir}+\frac{3a}{8ir}+\frac{1}{2r} & \leq & 1, \end{array}$$

or $r \geq 2 + \frac{3a}{8} + \frac{3a}{8} + \frac{1}{2} = 3a/4 + 5/2 = (3a + 10)/4$. □

The two conditions on r in Claims 33 and 34 are $r \geq 4a$ and $r \geq (3a + 10)/4$, respectively, where $a = \sqrt{\alpha c_1}$ while the condition on c_1 in Claim 33 is $c_1 \geq \frac{1}{\alpha} \exp(8 + \frac{128}{3\alpha^2})$. Since $c_1 \geq \frac{1}{\alpha}$, the former condition on r, $r \geq 4a$, is tighter. This leads to the following theorem.

Theorem 35. *For a fixed $\alpha > 0$, there is a fixed $\beta > 0$ such that in a market with average wealth $\Theta(r)$ that obeys the $(1 + 4\alpha)$-smoothness and the relaxed $(1 + 4\alpha)$-WGS properties, it is NP-hard to compute prices at which a $(1 - \beta/r)$ approximation is possible, where $r \geq 4\sqrt{\alpha c_1}$ and $c_1 \geq \frac{1}{\alpha} \exp(8 + \frac{128}{3\alpha^2})$.*

5 Relationship between discontent and ϵ-closeness in utility

In this section, we discuss the relationship between the two measures of distance from equilibria in the divisible setting. Specifically, for a fixed set of prices $\mathbf{p}$, let $\mathbf{x}^*$ denote the optimal allocation of a buyer b who has initial wealth w. We are interested in determining the discontent of a bundle $\mathbf{x}$ that is ϵ-close to $\mathbf{x}^*$ in terms of utility; i.e. $u(\mathbf{x}) = (1 - \epsilon)u(\mathbf{x}^*)$. Given a utility $\overline{u}$ (*not* a utility function), a set of prices $\mathbf{p}$, define $\psi(\overline{u}, \mathbf{p})$ as follows:

$$\psi(\overline{u}, \mathbf{p}) = \min\{\mathbf{p} \cdot \mathbf{x} \mid u(\mathbf{x}) \geq \overline{u}\}$$

In words, given a desired utility level $\overline{u}$ and set of prices $\mathbf{p}$, $\psi(\overline{u}, \mathbf{p})$ is the minimum amount of wealth required to afford a bundle with utility value $\overline{u}$. In the economics literature, $\psi(\cdot)$ is also known as the *expenditure function* [15].

The discontent for the bundle $\mathbf{x}$ is then given by $\psi(u(\mathbf{x}^*), \mathbf{p}) - \psi((1 - \epsilon)u(\mathbf{x}^*), \mathbf{p})$. In the divisible setting, it is fairly standard to assume that $\psi(u(\mathbf{x}^*), \mathbf{p}) = w$ [6] and thus the expression for the discontent reduces to $w - \psi((1 - \epsilon)u(\mathbf{x}^*), \mathbf{p})$. Define the relative discontent as

$$\epsilon_d = 1 - \frac{\psi((1 - \epsilon)u(\mathbf{x}^*), \mathbf{p})}{w}$$

Theorem 36. *If the underlying utility function $u(\cdot)$ is concave, then the relative discontent is no smaller than the ϵ-closeness in utility. That is $\epsilon_d \geq \epsilon$.*

Proof. It suffices to show that $\psi((1-\epsilon)u(\mathbf{x}^*), \mathbf{p}) \leq (1-\epsilon)w$. It is known that if the utility function $u(\cdot)$ is concave, then the expenditure function $\psi(\cdot)$ is convex in $\overline{u}$[15]. Thus, for any $\epsilon \in [0, 1]$, we have

$$\begin{aligned} \psi(\epsilon 0 + (1 - \epsilon)u(\mathbf{x}^*), \mathbf{p}) &\leq \epsilon\psi(0, \mathbf{p}) + (1 - \epsilon)\psi(u(\mathbf{x}^*), \mathbf{p}) \\ &= (1 - \epsilon)w \end{aligned}$$

□

For Cobb-Douglas utility functions, the two measures are identical. Let $u(\mathbf{x}) = \prod_i x_i^{\alpha_i}, \sum_i \alpha_i = 1$. Let $\mathbf{x}$ satisfy $u(\mathbf{x}) = (1 - \epsilon)u(\mathbf{x}^*)$, where $\mathbf{x}^*$ is the optimal allocation with cost w. By scaling the allocation $\mathbf{x}^*$ by a factor of $(1 - \epsilon)$, one achieves an allocation with utility $u((1 - \epsilon)\mathbf{x}^*) = (1 - \epsilon)u(\mathbf{x}^*) = u(\mathbf{x})$. But $(1 - \epsilon)\mathbf{x}^*$ is cheaper than $\mathbf{x}^*$ by ϵw units of money. Furthermore, $u(\mathbf{x})$ is

[6] If the buyer has utility for money, this is necessarily the case.

the largest utility one can achieve with a wealth of $(1-\epsilon)w$. The relative discontent is therefore ϵ, which is also what one would obtain if one took the ratio of the utility values.

To contrast our approach and the approach of using the ratio of utility values as the discontent, consider the utility function that is obtained by taking the log of the Cobb-Douglas utility function, i.e. $u_{\log}(\mathbf{x}) = \sum_i \alpha_i \log x_i$, $\sum_i \alpha_i = 1$. A utility functions of this form is called a log-linear utility function in the economics literature [15]. $u_{\log}(\cdot)$ carries identical preference ordering as $u(\cdot)$. Also, $u_{\log}(\cdot)$ is concave because for $x_i > 0, y_i > 0, t \in [0,1]$, we have:

$$\begin{aligned} t \log x_i + (1-t)\log y_i &\leq \log(tx_i + (1-t)y_i) \qquad \text{(from concavity of log)} \\ \Leftrightarrow \sum_i a_i(t \log x_i + (1-t)\log y_i) &\leq \sum_i a_i \log(tx_i + (1-t)y_i) \qquad \text{(since } a_i \geq 0) \\ \Leftrightarrow t u_{\log}(\mathbf{x}) + (1-t)u_{\log}(\mathbf{y}) &\leq u_{\log}(t\mathbf{x} + (1-t)\mathbf{y}) \end{aligned}$$

Suppose $u(\mathbf{x}) = (1-\epsilon)u(\mathbf{x}^*)$. As mentioned before, the discontent under $u(\cdot)$ is ϵ under both measures. Now let us examine what happens to the discontent, as measured by taking the ratio of the utility values, when one considers the utility function $u_{\log}(\cdot)$ instead of $u(\cdot)$. The discontent is:

$$\begin{aligned} 1 - \frac{u_{\log}(\mathbf{x})}{u_{\log}(\mathbf{x}^*)} &= 1 - \frac{\log u(\mathbf{x})}{\log u(\mathbf{x}^*)} \\ &= 1 - \frac{\log(1-\epsilon) + \log u(\mathbf{x}^*)}{\log u(\mathbf{x}^*)} \\ &= \frac{-\log(1-\epsilon)}{\log u(\mathbf{x}^*)} \\ &\geq \frac{\epsilon}{\log u(\mathbf{x}^*)} \end{aligned}$$

Thus, under the modified utility function $u_{\log}$ that carries the same preference information, the approach of measuring discontent by taking the ratio of actual utility to the optimal utility depends upon $\log u(\mathbf{x}^*)$ and can, in some cases, give an arbitrarily different measure of discontent.

References

[1] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

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

[3] Laurence Ausubel. Walrasian tatonnement for discrete goods. *Technical report, University of Maryland*, 2005. Available at http://www.cramton.umd.edu/workshop/papers/ausubel-walrasian-tatonnement.pdf.

[4] Bruno Codenotti, Benton McCune, and Kasturi Varadarajan. Market equilibrium via the excess demand function. In *Symposium on Theory of Computing (STOC)*, pages 74–83, 2005.

[5] Bruno Codenotti, Sriram Pemmaraju, and Kasturi Varadarajan. The computation of market equilibria. *SIGACT News*, 35(4):23–37, 2004.

[6] Bruno Codenotti, Sriram Pemmaraju, and Kasturi Varadarajan. On the polynomial time computation of equilibria for certain exchange economies. In *Symposium on Discrete Algorithms (SODA)*, pages 72–81, 2005.

[7] Bruno Codenotti, Amin Saberi, Kasturi Varadarajan, and Yinyu Ye. Leontief economies encode nonzero sum two-player games. In *Symposium on Discrete Algorithms (SODA)*, pages 659–667, 2006.

[8] Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra. On the complexity of price equilibria. *Journal of Computer System Sciences*, 67(2):311–324, 2003.

[9] Nikhil R. Devanur and Vijay V. Vazirani. The spending constraint model for market equilibrium: algorithmic, existence and uniqueness results. In *Symposium on Theory of Computing (STOC)*, pages 519–528, 2004.

[10] Rahul Garg and Sanjiv Kapoor. Auction algorithms for market equilibrium. In *Symposium on Theory of Computing (STOC)*, pages 511–518, 2004.

[11] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.

[12] Kamal Jain. A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In *Foundations of Computer Science (FOCS)*, pages 286–294, 2004.

[13] Kamal Jain, Mohammad Mahdian, and Amin Saberi. Approximating market equilibria. In *Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, 2003.

[14] Kamal Jain, Vijay V. Vazirani, and Yinyu Ye. Market equilibria for homothetic, quasi-concave utilities and economies of scale in production. In *Symposium on Discrete Algorithms (SODA)*, pages 63–71, 2005.

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

[16] Paul Milgrom and Bruno Strulovici. Concepts and properties of substitute goods. *Technical Report*, 2006. Available at http://ideas.repec.org/p/nuf/econwp/0602.html.

ECCC ISSN 1433-8092
http://eccc.hpi-web.de/