HOW GENERAL EQUILIBRIUM IN MARKETS WITH INDIVISIBLE GOODS OBTAINS AT THE COST OF COMPUTATIONAL COMPLEXITY

REPORT

PETER BOSSAERTS, KONSTANTINOS IOANNIDIS, ROBERT WOODS, NITIN YADAV

Contribution

- A long-standing question in general equilibrium theory is whether and how markets equilibrate in the presence of indivisibilities, especially under homogeneity of preferences and tight individual budgets.
- In that setting, everyone is actually solving an NP hard budget problem.
- Our conjecture is as follows: by increasing the complexity of agents' individual budget problems, markets "scatter" otherwise homogeneous individuals in terms of their willingness or capacity to spend cognitive effort to find good consumption bundles; the resulting heterogeneity in cognitive effort potentially restores equilibrium existence.
- In a laboratory market experiments with 3 goods and money, market prices indeed push individual budget problems into the region of most difficult instances; this happens in spite of our attempts to make budget allocations simple by charging suppliers zero marginal cost.
- Three approximately equal-sized groups of consumers emerge, stratified by how difficult it was to implement the choices they made; the group that made the most difficult choices end up earning most.
- When it exists, the Walrasian equilibrium is different, but fails to emerge.

Date: December 2024.

Bossaerts, Ioannidis and Woods are from the University of Cambridge, Yadav is from the University of Melbourne. The research reported here was generously funded by the Leverhulme Trust. We are grateful to conference participants at the 14$^{\text{th}}$ Society for Experimental Finance (Stavanger, 2024) and the 1$^{\text{st}}$ Annual Conference of the Network of the UK-Based Experimental and Behavioural Economists (Leicester, 2024), and seminar attendees at Royal Holloway University of London, the University of Exeter, and the University of Vienna for their comments.

Contents

1. Introduction

Example of equilibrium non-existence:

Take multiples of 3 identical consumers with income (cash) equal to £2.80 who compete to obtain:

- *One trip to London (L), valued at £1.20, and/or*
- *One MacBook Pro (M), valued at £2.40, and/or*
- *One Hermès bag (H), valued at £3.00.*

Utility is additive (including cash remaining).
Supply of these goods: 2 (per 3 consumers).

Question: How can prices be set so that demand = supply?

Answer: Each agent needs to be indifferent between pairs of goods and refrain from buying a third good. But this is impossible, unless income is sufficiently large (here, at least £3.20: with tick size of 5p, equilibrium prices are then equal to 0.10 (L), 1.30 (M) and 1.90 (H)).

But this ignores difficulty to find the optimum solution. The traditional Walrasian equilibrium assumes that everyone can find the optimum budget allocation given prices. This may not be true as, under indivisibilities, the budget allocation problem is "NP hard." If markets choose prices such that the instance of the budget allocation problem agents face is most complex, then heterogeneity in cognitive effort applied may lead different agents to choose different goods baskets, allowing the market to equilibrate. Importantly, prices have to be set so that the cognitive effort is compensated for: as effort to reach solutions increases, the corresponding utility gained has to increase.

2. NP-Hard Problems: A Digression

We will argue that the budget allocation problem will have to become "hard" for markets to equilibrate. Here, we discuss the meaning of hardness. Details can be found in standard textbooks such as Arora & Barak (2010).

To fix ideas, below are examples of "NP hard problems", namely: 0-1 Knapsack Problem (KP, Figure 1a), and Traveling Salesperson Problem (TSP, Figure 1b).

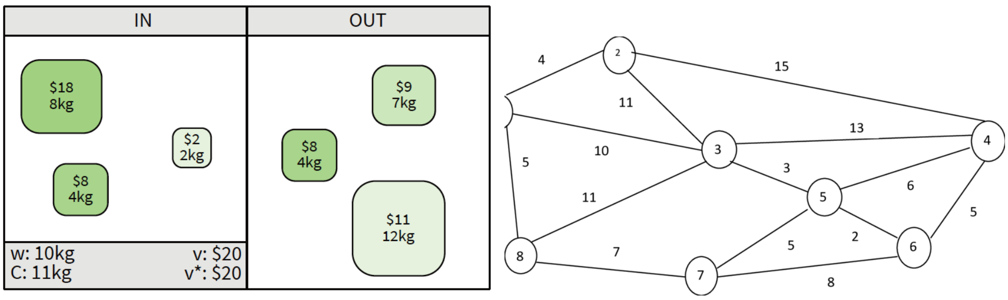

(A) 0-1 Knapsack Problem (B) Traveling Salesperson Problem

Left: 0-1 Knapsack Problem; an NP hard problem where one must find the highest possible sum of values of items to include 'in' the knapsack or reach a benchmark value v^* without the sum of weights (w) of the items going over the capacity (C) of the knapsack. Each item has its own value and weight and can be dragged 'IN' or 'OUT' of the knapsack. Items are indivisible. v: value of the knapsack based on items currently 'IN.' Right: Traveling Salesperson Problem; an NP hard problem where one must visit every node (city) once and only once in the shortest distance (least cost). Each edge is a travel path and can be selected to travel in either direction.

Figure 1. Examples of NP-hard problems

NP stands for "Non-deterministic Polynomial." It will be distracting to explain this term, since it would get us into exploring non-deterministic Turing machines, a generic from of computing device. Instead, we present the theory here without reference to Turing machines (deterministic or non-deterministic).

2.1. **The class NP.** We consider two types of problems, namely:

- NP *decision* problems (yes/no). E.g.: Can one reach a knapsack value of $v^* = \$20$ in Figure 1a. Is it possible to find a path that connects all nodes and visits each node only once while the total length of the path is at most, say, 70 in Figure 1b? The decision problems belong to the class "NP" if it is easy to verify a solution (in time polynomial in its size N). But it may be hard to find it.
  - Notice the asymmetry: if the solution = "no," it may be really hard to prove (co-NP).
- NP *optimization* problems. E.g.: what is the best knapsack one can reach in Figure 1a? What is the shortest path that connects all nodes in Figure 1b? Notice that optimization problems are sequences of decision problems with increasing criterion.
  - So, verifying that one has reached an optimum is itself an NP problem, and hence, could be really hard.
  - Contrast this with ordinary, optimization problems with $\mathcal{C}^2$ functions (twice-differentiable), where one can use first and second derivatives to check whether one has reached a optimum.

2.2. **The Class NP complete (and hard).** We will focus on the equivalence class "NP complete," decision problems that all other problems in the class can be reduced to within a finite number of steps:

- If an instance of size N of, say, the TSP is solvable (answer = "yes") then the corresponding instance of the KP of size N is also solvable, and *vice versa*.
- The number of steps in the reduction must not increase more than a polynomial in N.

How hard are these? We do not know. At present, there does not exist an efficient (polynomial-time) algorithm to solve them.[1]

Consider them to be an interesting class of "equivalent" problems; humans appear to deal with them in distinct ways (see later). They are not probabilistic tasks like a multi-armed bandit problem.

[1]The problem is referred to as the question whether "P = NP."

The class "NP hard" also includes optimization problems (sequence of decision problems).

2.3. **The consumer budget problem with indivisibilities is a KP, and hence, NP hard.** (Gilboa et al. 2021)

- Goods: $i = 1, \dots, I$; with payoffs π_i; and prices ("costs") p_i.
- Cash (income): C.

The agent faces the following optimization problem.

$$\max_{x_i \in \{0,1\}, i=1,2,\dots I} \sum_i x_i \pi_i + \left(C - \sum_i x_i p_i \right)$$
$$\text{subject to:} \sum_i x_i p_i \leq C.$$

The objective function (utility) can be rewritten as follows:

$$\max_{x_i \in \{0,1\}, i=1,2,\dots I} C + \sum_i x_i (\pi_i - p_i).$$

Because C is constant, we have a standard 0-1 KP problem with values $(\pi_i - p_i)$ and costs p_i.

2.4. **Algorithm-dependent metrics of instance complexity of KP optimization.** As we propose that markets will select for price configurations that make agents' budget problems sufficiently complex for markets to equilibrate, we need a measure of instance complexity, i.e., complexity of a given KP with given values and prices.

Here is one that has proved to be successful in predicting performance and effort of humans, at both the individual (Murawski & Bossaerts 2016) and market levels (Meloso et al. 2009, Bossaerts et al. 2024). This one assumes that agents follow a particular algorithm in solving the KP. It is referred to as the Sahni-Horowitz algorithm; see Meloso et al. (2009).[2]

The metric starts with the greedy algorithm, which prescribes filling the knapsack with items in order of descending ratio value/weight till capacity is reached. If the pivotal item could be split in pieces to fit the knapsack, then one would

[2]Notice that the Sahni-Horowitz algorithm guarantees finding the optimal solution, while humans generally only approximate the optimal solution. The metric, Sahni-k, provides a way to understand why.

obtain the optimal knapsack under infinite divisibility. This shows how simple the traditional budget problem (with infinitely divisible goods) is. There are two versions of the greedy algorithm. The "amended" version continues attempting to add further items by bypassing items which do not fit in the remain knapsack capacity, rather than just stopping at the first item that can no longer be fit.[3]

Sahni-k =

0: *If greedy algorithm finds the solution*
1: *If one item has to be put in knapsack before the greedy algorithm can find the solution*
2: *If two items have to be put in the knapsack before the greedy algorithm can find the solution*

...

2.5. **Algorithm-independent metrics of instance complexity of KP optimization.** It would be better if we could use a metric of instance complexity that does not make an assumption as to how humans solve (or attempt to approximate solutions of) the KP since human attempts appear to be extremely heterogeneous – both across individuals and time; see Murawski & Bossaerts (2016).

In analogy with other NP hard problems (the TSP and the 3-SAT problem), we proved that there exists a "phase transition" in the space of normalized capacity κ and normalized profit u where effort required is highest, though only for algorithms that find the solution. See Yadav et al. (2020*b*). In our setting:

$$\kappa = \frac{C}{\sum_i p_i},$$
$$u = \frac{U}{\sum_i (\pi_i - p_i) + C}.$$

Here, U is the target utility: $U = \sum_i x_i^*(\pi_i - p_i) + C$, where x_i^* ($\in \{0,1\}$) is the optimal choice.

[3]The amended version has good approximation properties (50% of the optimal value is guaranteed) provided one adds one more check: whether the value attained is no smaller than the value of the most profitable item that can fill the knapsack on its own. If not, then this latter item should be chosen instead. In computing Sahni-k, we ignore this complication.

Phase Transitions: The most difficult instances occur where $u - \epsilon < \kappa \leq u$ for some small ϵ. The upper boundary is convex.

Intuition:

- Low values of κ imply that the instance is over-constrained. Higher values allow for more feasible allocations ("feasibility")
- High values of u imply that the utility is hard to reach; even if many choices are feasible, there may only be a few choices that get close to u ("solvability" can be obtained only for very few values close to u)
- The most difficult instances are where feasibility and solvability are critical: lots of feasibility but hardly solvable, or hardly feasible while easy to increase utility (incidentally, just determining which choices are feasible requires effort!)

Example: prices are such that one can afford almost everything – but not exactly (high κ), while the target u is high (u requires almost almost all goods).

Figure 2 is taken from Yadav et al. (2020*b*). On top is the theoretical prediction (Figure 2a). The red line is a specific instance we will talk about later. At the bottom are times spent before choosing to either abandon further search for improvement (declaring the chosen knapsack to be optimal) or moving to a better knapsack, when the individual (Figure 2b) or computer (Figure 2c) had reached a knapsack with indicated log-ratio of κ and u. NOTE: In these plots, notation is different: $c = \kappa$ and $u = p$. The computer used general-purpose optimization software, namely, Minizinc.

Notice substantial effort spent by both humans and computers on instances where $\log(c/p) \equiv \log(\kappa/u) \ll 0$. At that point, the instance is over-constrained, and hence, "unlikely" to be solvable, so the agent could as well guess that they are not solvable and stop search. Of course, computers do not have a choice: they have to work until convinced that there are no further improvements possible. Both humans and computers spent "irrational" amounts of efforts to find improvements. We will find that this biases prices in markets: they will make it almost impossible for agents in the marketplace to discover how to enhance utility through trade.

2.6. More on individual human choice:

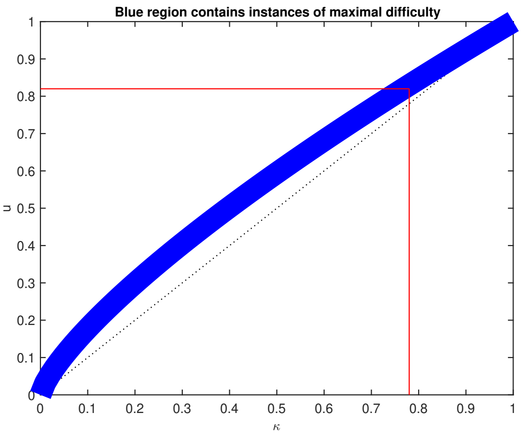

(A) Location of phase transition in κ-u space

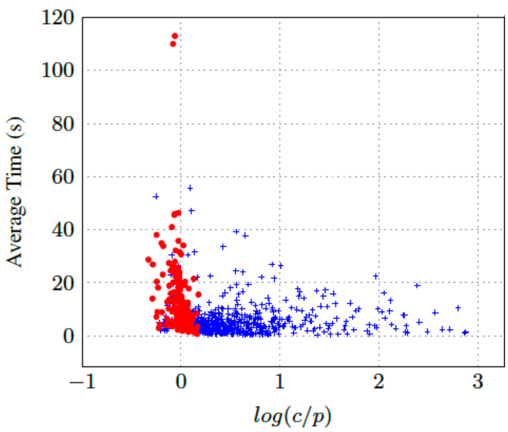

(B) Time spend by humans

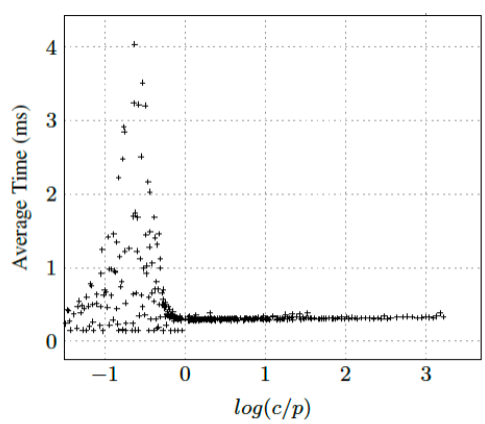

(C) Time spend by computers

Top: A: Plot of theoretical area of most difficult instance (blue) in $\kappa - u$ space. B, C: Observed times in knapsack problems for humans (B; N=10 items) and computers (C; N=20 items) to augment value of chosen knapsack or declare that an optimum has been reached, as a function of $\log(\kappa - u)$; in the bottom figure, $c = \kappa$, $p = u$. Reproduced from Yadav et al. (2020*a*). For humans, the red nodes are local maxima, which means that the knapsack was full and hence, to improve knapsack value, the individual had to *replace* at least one item.

FIGURE 2. Theoretical region of maximal complexity (A) and empirical record of solution times (B, C)

- Unlike in, say, multi-armed bandit problems, humans (and non-human primates) *do NOT use trial-and-error*; they perform substantially better (Murawski & Bossaerts 2016).
- Humans do not necessarily find the optimum even for small problems; they *approximate* (Murawski & Bossaerts 2016).

- Performance and effort correlate with *approximation complexity classes* (Yadav et al. 2022).
- Humans *search in heterogeneous ways* (across people; across time) (discussed before) (Murawski & Bossaerts 2016).
- *Performance can be modeled as a function of instance complexity*, e.g., Sahni-k (Meloso et al. 2009, Murawski & Bossaerts 2016, Franco et al. 2021).[4]
- *Effort can be modeled as a function of instance complexity*, e.g., Sahni-k (Murawski & Bossaerts 2016, Franco et al. 2021, 2024).[5]

3. How markets could equilibrate by exploiting complexity

Reasoning:

- Budget problems are computationally hard ("NP hard"); they are a kind of 0-1 knapsack problem (KP)
- Humans are very heterogeneous in the way they solve (or approximate solutions of) KPs – across individuals and over time
- **Heterogeneity could be exploited if markets set prices in order to make the budget problem "hard," and hence, split demand:**
  - Those who spend most effort find the optimal solution and gain most (buying, say, H & M)
  - Those who spend medium effort find a good solution so gain a bit less (buying, say, H & L)
  - Those who spend least effort end up with the smallest gains (buying, say, M & L)
- By splitting demand, an equilibrium could obtain whereby cognitive effort is rewarded – unlike in the (existing or non-existing) Walrasian equilibrium

[4]A meta-analysis of 6 studies with in total 4705 participant-trials shows that increasing Sahni-k by one notch decreases value attained (as percentage of optimal) as much as increasing the number of available items by 5. There is significant interaction, though: increasing the number of items by 1 reduces the effect of Sahni-k by as much as the effect of the increase in number of items. *Unpublished Statistics.*

[5]A meta-analysis of 6 studies with in total 4705 participant-trials shows that increasing Sahni-k by one notch increases time spent as much as increasing the number of available items by 1. *Unpublished statistics.*

The new definition of equilibrium extends the Walrasian equilibrium by requiring effort to be compensated when prices are such that the optimal budget allocation is difficult. Even if the Walrasian equilibrium exists, it does not compensate for cognitive effort needed to reach the optimum even if there exists multiple optima, and hence, the difficulty of at least one optimum has to be higher than that of another.

Numerical example based on algorithm-dependent instance complexity. With a budget capacity of 2.80, the following is an equilibrium:

Good	Payoff	**Price***	Ratio Price/Payoff λ
L	1.20	0.40	1/3
M	2.40	1.20	1/2
H	3.00	1.60	8/15

*Other prices can support this equilibrium

Consumer	Algorithm		**Choice**	Earnings
	Applied	Solves OK If		
1	Greedy	(Sahni-)$k=0$	(M,L)	4.80
2	(1 Item +) Greedy	$k=1$	(H,L)	5.00
3	(2 Items +) Greedy	$k=2$	(H,M)	5.40

Discussion: Our predictions may seem imprecise, but remember that economic theory has yet to produce ANY prediction (when Walrasian equilibrium does not exist). We can imagine some *ad hoc*, half-baked, alternative theories:

- Prices are fixed at a level that all goods are affordable but demand is rationed (because of our experimental design, prices will be very low); the key is to trade quickly; trade will subside fast
- Prices are volatile and unpredictable (within and across replications) so some consumers end up with no goods, others with 1 good, still others with 2 goods, and remainder end up with all three goods; no particular prediction as to how many in each cohort: allocations happen by chance

When the Walrasian equilibrium does exist, it is not obvious that markets will go for it, since the Walrasian equilibrium does not compensate for difficulty, while with indivisibilities, the budget allocation problem generically has to be complex ("NP hard").

Our experiment will verify what happens both when the Walrasian equilibrium exists and when it does not exist.

4. The experiment: Design

In each experiment session, 16 or 20 participants trade 3 goods for cash in an online marketplace for short periods of time (3 minutes), after which markets close and earnings are computed based on remaining cash and liquidating values of the final holdings of the goods. Trade takes place on Flex-E-Markets, an online, multi-market electronic continuous open-book system based on limit orders and with immediate settlement. Snapshot of interface can be found in Figure 3.

Flex-E-Markets is an online, multi-market electronic continuous open-book system based on limit orders and with immediate settlement. Panels, in clockwise order: Holdings, Order Book and Trade History (Good A, B and C), Order Form. Holdings presents traders with their current balance of goods and cash. The Order Book presents a anonymous record of all standing orders with the price and quantity (but not the traders); trade history is displayed in the "Trade History" column. Cancellation of a trader's own standing orders can be done directly from the Order Book panel. The Order Form is for submission of new orders for the chosen quantity and limit price. Goods are referred to as A, B and C rather than L, M and H. Currency sign is \$; \$1 = £1 (currency of payment).

Figure 3. Flex-E-Markets Trading interface

The Summary of the task for participants (from the Instructions) is re-produced in Figure 4.

Trading Game

In the trading game, in each period you will be assigned the role of either a **Seller** or a **Consumer**. You can find out which you are by looking at your holdings. This is important, so please note it each period.

A **Seller** will start with a positive quantity of the assets, but no cash. A Seller gets **NO** value from holding assets, but they do for cash, so they want to sell their assets for as much cash as possible. Any assets not sold by a Seller at the end of the period will be discarded for no value.

A **Consumer** starts with cash, but no assets. Assets provide significant value to the Consumers, so Consumers will likely want to use their cash to purchase assets. Importantly, Consumers only get value from the **FIRST** unit of each asset they hold, plus any remaining cash they have left over. A Consumer holding two or more units of an asset will not collect value for the subsequent units; only the first unit will count.

Figure 4. Summary of Experiment Instructions

Notice that *supplies* of goods are in the hands of suppliers (Sellers in the instructions) with zero marginal cost. Consumers, assigned homogeneous values for the goods, start only with cash (income). This maximizes trade, since all goods need to move from suppliers to consumers. We want to observe as much trade prices as possible, to maximize power.[6] There is an additional, important consideration to this design:

- To ensure equilibration, prices have to be **high**: everyone can afford pairs, but not triplets; κ should be high.
- But suppliers have zero marginal cost, and this is known to push prices – under divisible goods – to **very low** levels, way below Walrasian equilibrium; see Smith & Williams (1982); recently replicated in Rasooly (2022).
  - There are Marshallian explanations for this; see Asparouhova et al. (2024).

[6]Contrast this with a design such as in Smith et al. (1988): assuming risk neutrality, we do not expect a single trade, so we have no observations – in expectation – with which to verify whether prices are right. Analogous problems exist for some of the designs in Plott & Sunder (1982, 1988).

We therefore exploit Marshallian dynamics to stack the deck against our complexity-based equilibrium theory.

Treatments:

1. *Equilibrium Existence:* Two types of 5 (five) sessions each (total: **10 sessions**):
   N: Walrasian equilibrium does not exist because income (cash) of consumers is too low.
   E: Unique Walrasian equilibrium exists because income (cash) of consumers is sufficiently high.
2. *First-Good Values*: Two goods-values configurations over 8 repetitions ("periods") each (total: **16 periods per session**):

	First-Good Values		
Configuration	**L**	**M**	**H**
I	£1.20	£2.40	£3.00
II	£1.60	£2.40	£3.20

Each configuration was run for 8 successive periods in a session. The order the configurations were run in was alternated across session. A consumer holding 1 unit of good H under goods-values configuration I at the end of the period would be paid £3.00 plus the value of their other goods and cash if that period was selected as a payment period. If a consumer held two or more units of any good, they would only be paid for the first unit of the good. Suppliers received nothing for the goods they held at the end of the period.

Table 1. First-good values for configurations I and II

We set tick size equal to 0.05 so the Walrasian equilibrium in Treatment E is unique, and price of L equals 0.10. (Trade at price of zero is not allowed and trading at 5 cents would afford purchase of all three goods, inconsistent with equilibrium.) Participants rotate roles (supplier: 1 out of 4 periods; consumer: 3 out of 4 periods); they are paid randomly for three periods where they were consumer, and one period where they were supplier.

5. The experiment: Predictions

Expected outcomes:

(1) Prices will be high, so that any 2 but not all 3 goods can be afforded

(2) Prices will be even higher for Treatment E, in contrast with Walrasian equilibrium, where price of lowest-value good equals minimal possible (0.10)
(3) Efficiency will be high: most goods are sold to consumers (90%+?)
(4) The vast majority of consumers end up holding goods in pairs
(5) Approximately the same number of consumers end up holding each (of 3) pairs of goods
(6) Prices of goods tend to configurations where the budget problem instance for consumers in the region of phase transition
(7) When the location of the budget problem instance is outside the phase transition, it tends to be where $\log(\kappa/u) \ll 0$
(8) Consumers are expected to gain most utility when buying the pair of goods most difficult to reach, and to gain least when buying the pair least difficult to reach.

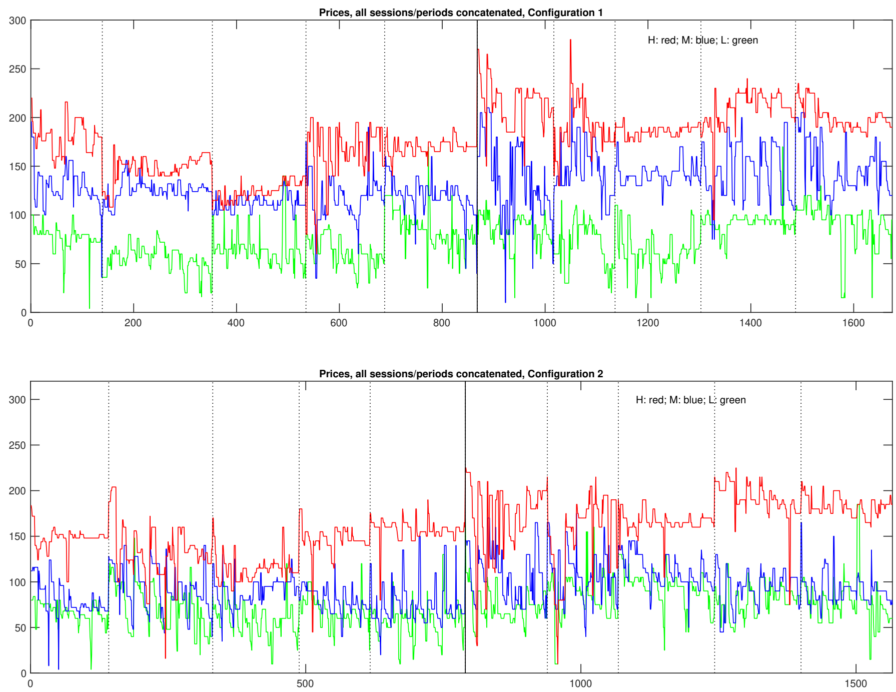

Top: First-Good Values Configuration I; Bottom: First-Good Values Configuration II. Dotted vertical lines delineate sessions. Solid vertical line indicates whether Walrasian equilibrium exists: to the left, Walrasian equilibrium does NOT exist (N Treatment); to the right, it exists (E Treatment). Good colors: H - red, M - blue, L - green.

FIGURE 5. Time-series of goods prices by configuration

6. The experiment: Results

Here is the evidence.

1. Prices will be high, so that any 2 but not all 3 goods can be afforded.
Figure 5 depicts the evolution of trade prices. Prices are high, contrary to the case with divisible goods and low marginal cost for suppliers (Smith & Williams 1982). Budgets are constrained but pairs can be afforded (κ is high). Prices are volatile though, making it hard for consumers to decide what to buy. This adds to complexity.

2. Prices will be even higher for Treatment E, in contrast with Walrasian equilibrium, where price of lowest-value good equals minimal possible (0.10).
Figure 5 shows that, when Walrasian equilibrium exists (to Right of solid vertical line; Treatment E), prices are even higher than when it does not. Under Walrasian equilibrium, price of item L (green) should be 10 (cents). Prices are uniformly above it; in fact they are between five and ten-fold too high.

3. Efficiency will be high: most goods are sold to consumers.
Table 2 proves that most suppliers sell most of their holdings. Efficiency is 90% on average (90% of the goods move from suppliers to consumers). Very few suppliers are left with 1 up to 5 units of the goods (they are endowed with 6 units to start with; 2 of each good).

4. The vast majority of consumers end up holding goods in pairs.
Table 3 shows that $\approx 2/3$ (68.6%) of consumers end up with a pair of the goods, consistent with equilibrium. Still, 1/5 (19.6%) manages to buy only one good. Remarkably, about 1/20 consumers (4.5%) ends up holding more than 1 unit of a good, suggesting that speculation is significant (buy low to sell high). Those consumers will not be paid a liquidating dividend for their additional units of holdings (consumers are only paid for the first unit of each good they hold) and it keeps other consumers from increasing their utility. As a result, the efficiency numbers computed before over-estimate true efficiency. Since consumers speculate with about three percent (2.8% to be precise) of the total supply of goods, average

TABLE 2. Supplier End-of-Period Holdings

Treatment	**N (No WE)**	**E (WE)**	**Combined**
Supplier-periods	352	368	720
Fraction Endowment Sold	91.7%	88.1%	89.9%
	(1936 / 2112)	(1946 / 2208)	(3882 / 4320)
Remaining goods	None: 263	None: 256	-
	1: 41	1: 49	
	2: 26	2: 22	
	3: 12	3: 18	
	4: 3	4: 8	
	5: 7	5: 10	
		6: 5	

The supplier holdings of goods at the end of all periods.
Supplier-periods: the total number of periods multiplied by the number of suppliers in each period.
Fraction Endowment Sold: Percentage of goods endowed with that suppliers sell by the end of the period.
Remaining goods: Tally of supplier-periods where the number of goods indicated (None, 1, 2, ...) were not sold. "None" means suppliers sold all their goods in a period. "4" means they sold 2 units (of their starting 6) and had 4 remaining.

efficiency is not 90%, but about 3 percentage points lower, and stands at 87%. Speculation is encouraged by the observed volatility in prices.[7]

5. Approximately the same number of consumers end up holding each (of 3) pairs of goods.

For equilibrium to obtain, consumers need to end up with pairs of goods in equal proportion. Table 4 shows that approximately the same number of consumers (consumer-periods) end up holding the combinations ML, HL and HM, suggesting that equilibrium allocations are reached. We do observe a slight bias towards purchases of the highest-value good, H. (For reference: total consumer-periods number = 2160.)

[7] "No WE" refers to Treatment N; "WE" refers to Treatment E (when Walrasian equilibrium exists). Number of consumer-periods equals number of consumers multiplied by number of periods across all sessions of Treatment N or E.

Table 3. Consumer End-of-Period Holdings

Treatment	N (No WE)	E (WE)	Combined
Consumer-periods	1056	1104	2160
All (HML)	4.8%	2.1%	3.4%
Pairs (HM, HL, or ML)	70.0%	67.2%	68.6%
Singles (H, M, or L)	16.9%	22.2%	19.6%
None (-)	4.0%	3.7%	3.8%
Excess (E.g., HLL)	4.2%	4.8%	4.5%
Excess units count of supply	2.5%	2.9%	2.8%
	(52 / 2112)	(64 / 2208)	(116 / 4320)

Consumer-periods: the total number of periods multiplied by the number of consumers in each period. *All*: percent of consumers who ended the period holding 1 unit of all 3 goods (H, M, and L). *Pairs*: percent of consumers who ended the period holding 1 unit of 2 goods but not the third (H and M, H and L, or M and L). *Singles*: percent of consumers who ended the period holding 1 unit of 1 goods but not the other 2 (H, M, or L). *None*: percent of consumers whp ended the period holding 0 goods. *Excess*: percent of consumers who ended the period holding more than 1 unit of any of the assets (e.g., H, L, and L). *Excess units count of supply*: percent of goods units which were held in excess (more than 1 unit held by a single consumer), and as such did not receive compensation.

6. Prices of goods tend to configurations where the budget problem instance for consumers in the region of phase transition.

The plot in Figure 6 shows the locations, in $\kappa - u$ space, of the budget problem instances for each trade price configuration (i.e., after each trade in a good, keeping the prices of the other goods at the level of their last trade), stratified per session and first-goods values configuration where 1 is for first-goods values configuration I and 2 is for first-goods values configuration II. E.g., 6.1 contains all trades across all periods in session 6 for first-goods values configuration I; in Sessions 1-5, the Walrasian equilibrium did not exist (Treatment N); in Sessions 6-10, the Walrasian equilibrium existed (Treatment E). As expected, the capacities of the budget problem instances are not too constrained (κs are high); most often, consumers could afford 2 of the 3 goods, though when κ dropped to ≈ 0.5, consumers could only afford 1 of the 3 goods: capacity became constrained. Most locations are above the 45 degree line, but close to it. This is where the most

Table 4. Consumer End-of-Period Holdings By Goods

Treatment		**N (No WE)**	**E (WE)**	**Combined**
All	HML	51	23	74
Pairs	ML	209	254	463
	HL	257	266	523
	HM	273	222	495
Single	L	40	33	73
	M	62	65	127
	H	77	147	224
None	-	43	41	84
Holds good	L*	557	576	1133
	M*	595	564	1159
	H*	658	658	1316

Shown are consumer holdings of goods at the end of each period broken down to the exact goods they held. Any goods held in excess ($> 1unit$) were ignored. *Holds good*: the number of consumers which held at least 1 unit of the given good at the end of the period; for example; L* means the consumer held at least 1 unit of good L at the end of the period, regardless of all other goods held.

difficult instances reside. It means that cognitive effort had to be applied in order to find the best combination. Cognitive effort required was further reinforced because prices were volatile, quickly invalidating the appropriateness of a choice when prices changed.

While prices were volatile, implying that location in $\kappa - u$ space changed rapidly, we discovered a strong tendency for location to revert to just above the 45 degree line, where the budget problem instances are more difficult. We verified this by running an autoregression on the evolution of $\log(\kappa/u)$ and test whether the (long-run) expectation is in the small space of most difficult budget problem instances: $-\epsilon < \log(\kappa/u) < 0$ for some small $\epsilon > 0$. The number of lags in the autoregression was set sufficiently high so that its error terms had insignificant autocorrelations. The results of the autoregression are displayed in Table 5.

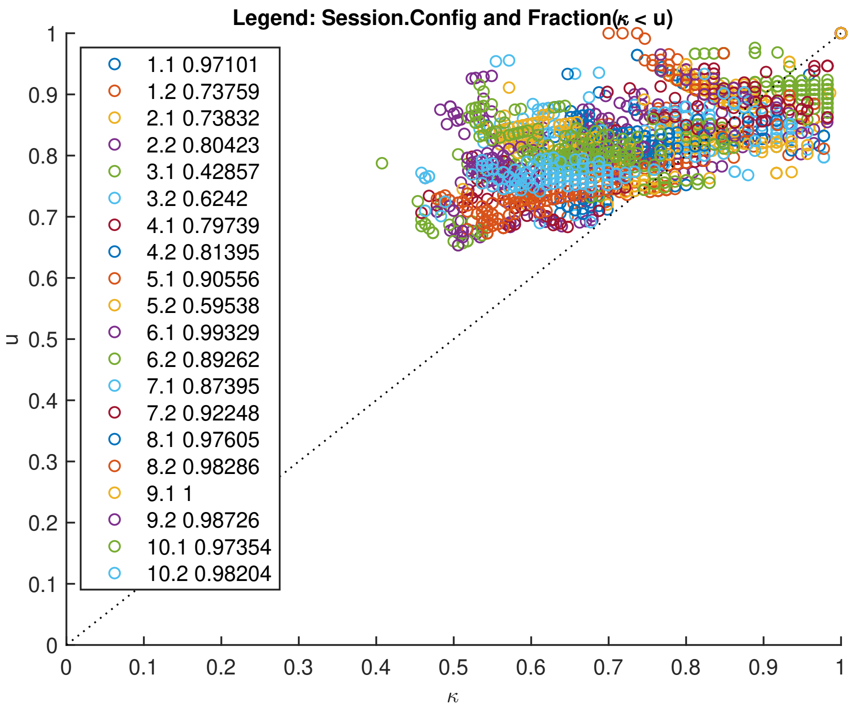

Each observation represents a trade in a good, keeping the prices of the other goods at the level of their last trade. Results are stratified per session and per goods-values configuration. The first number is the session number and the second number is the goods-values config where 1 is for goods-values configuration I and 2 is for goods-values configuration II. E.g., 6.1 contains all trades across all periods in session 6 for goods-values config I. Treatment N (the Walrasian equilibrium did not exist) is sessions 1-5; Treatment E (the Walrasian equilibrium did exist) is sessions 6-10.

FIGURE 6. Budget Problem Instances in $\kappa - u$.

One autoregression was run per session–configuration, concatenating all periods within a session–configuration.[8] Estimated long-term means of $\log(\kappa/u)$ are uniformly non-positive but small . Estimated long-term means tend to be farther below zero in Sessions 6-10, where the Walrasian equilibrium exists. This means that, given normalized capacity constraint (κ), the normalized profits (U) that could be made optimally were relatively higher given available net profits (differences between good value and good price), i.e., u was much higher. In other words, fewer feasible allocations reached utilities as high as the maximal utility. This is the effect of the increased goods prices we observed in Treatment E.

[8]Such concatenation is warranted because, as in prior markets experiments, prices in a replication usually start at the same level as they finished in the previous replication.

Table 5. Estimated Long-term Values of $\log(\kappa/u)$

Session:Config	$\log(\kappa/u)$	Session:Config	$\log(\kappa/u)$
1:I	-0.12^{*}	6:I	-0.24^{*}
1:II	-0.05^{*}	6:II	-0.19^{***}
2:I	-0.03	7:I	-0.18
2:II	-0.08^{*}	7:II	-0.17^{**}
3:I	0.00	8:I	-0.12^{***}
3:II	-0.01	8:II	-0.19^{***}
4:I	-0.05^{*}	9:I	-0.27^{**}
4:II	-0.05^{*}	9:II	-0.24^{***}
5:I	-0.11^{***}	10:I	-0.21^{*}
5:II	-0.03^{**}	10:II	-0.20^{***}

The long-term mean of $\log(\kappa/u)$ split by existence treatment (Left: N, i.e., Sessions 1-5; Right: E, i.e., Sessions 6-10), first-good values configuration (I or II) and session (1 to 10).
$^{***}: p \leq 0.001$; $^{**}: p \leq 0.01$; $^{*}: p \leq 0.05$.

In Table 5, significance levels are taken to be those of the intercept of the autoregression. Autoregression coefficients of low order were invariably extremely significant.[9]

7. When the location of the budget problem instance is outside the phase transition, it tends to be where $\log(\kappa/u) \ll 0$.

The many observations where the location of u is above κ correspond to situations when the budget problem instance where it was unlikely to find the optimum: while feasible allocations were reduced (κ is reduced), in fewer cases could one obtain utility improvements as large as the maximally obtainable ones. In other words, it was not easy to identify which trades were the best. We observed before that humans tend to spend a lot of time on such instances. In the context of markets, this means that prices may not change quickly, and the location of the

[9] In a future version, we will report significance levels of the estimated long-term average level of $\log(\kappa/u)$.

budget problem instance may remain a long time way above the 45 degree line in $\kappa - u$ space, consistent with the observations. The plot in Figure 6 shows, per session-configuration, the fraction of observations where $\log(\kappa/u) < 0$ ($\kappa < u$). With only one exception, the fractions are always above 0.5; in one case, it is even equal to 1 (Session 9, First-Goods Values Configuration I).

8. Consumers are expected to gain most utility when buying the pair of goods most difficult to reach, and to gain least when buying the pair least difficult to reach.

The pair of goods most difficult to reach in equilibrium is the pair (H, M): given income, these two goods tend to be the most expensive. With our new notion of equilibrium, participants who manage to choose this pair should be compensated for the extra cognitive effort required, in the form of higher increase in utility. Those who buy the cheapest pair (M, L) should be compensated least, since it is easiest to obtain. This is in sharp contrast with the Walrasian equilibrium, where any equilibrium choice (any pair of good) should lead to equal utility increase – agents are to be indifferent between those choices – since cognitive effort is ignored.

We can test whether prices tend to compensate purchasers of (H, M) most while purchasers of (M, L) receive least utility improvements by studying the long-term mean of the following vector autoregression.

Let $v_{(x,y)}(t)$ denote the earnings as of trade t from buying the pair (x, y). We run a vector autoregression on the difference in earnings between (H, L) and (M, L) and the difference in earnings between (H, M) and (H, L). Each observation corresponds to a trade, at which point we can re-evaluate these earnings differences (we keep the prices of the goods that did not trade equal to their most recent trade prices).

$$\begin{bmatrix} v_{(L,H)}(t) - v_{(L,M)}(t) \\ v_{(H,M)}(t) - v_{(L,H)}(t) \end{bmatrix} = \mu + \Phi_1 \begin{bmatrix} v_{(L,H)}(t-1) - v_{(L,M)}(t-1) \\ v_{(H,M)}(t-1) - v_{(L,H)}(t-1) \end{bmatrix} + \cdots + \Phi_K \begin{bmatrix} v_{(L,H)}(t-K) - v_{(L,M)}(t-K) \\ v_{(H,M)}(t-K) - v_{(L,H)}(t-K) \end{bmatrix} + \epsilon_t.$$

We test whether the long run expectation of the difference in values is strictly positive, i.e., whether

$$(I - \Phi_1 - \cdots - \Phi_K)^{-1}\mu$$

is strictly positive. The test can be done on the basis of concatenation of periods per session-configuration, as we did when estimating the long term mean of $\log(\kappa/u)$. Table 6 displays the estimation results. Significance levels are those of the intercept μ; first-order autoregression coefficients display sufficiently high significance levels that their estimation error can be ignored in a first passage at the test.

Table 6. Estimated Long-term Pairs Values

Session :Config	Long Term $v_{(L,H)} - v_{(L,M)}$	Long Term $v_{(H,M)} - v_{(L,H)}$	Session :Config	Long Term $v_{(L,H)} - v_{(L,M)}$	Long Term $v_{(H,M)}$ $v_{(L,H)}$
1:I	6	71**	6:I	−20	70*
1:II	13	66***	6:II	13	46***
2:I	38***	49	7:I	14	43**
2:II	41	58***	7:II	30	58***
3:I	47*	73***	8:I	15	42***
3:II	57***	48*	8:II	23	62***
4:I	7	65***	9:I	−4	66**
4:II	15*	67***	9:II	−18*	68***
5:I	8*	84***	10:I	6	66*
5:II	−2*	63***	10:II	−13*	66***

The long-term mean of the pairs values split by existence treatment (Left: N, i.e., Sessions 1-5; Right: E, i.e., Sessions 6-10), first-good values configuration (I and II) and session (1 to 10).
$v_{(L,H)} - v_{(L,M)}$ is the difference in value between the pair L and H and L and M.
$v_{(H,M)} - v_{(L,H)}$ is the difference in value between the pair H and M and L and H.
*** : $p \leq 0.001$; ** : $p \leq 0.01$; * : $p \leq 0.05$

The results confirm that cognitive effort is compensated for. In all but one cases, estimated earnings from (H, M) are significantly higher than those for the next-best pair (H, L), and mostly sizeably so (an estimate of 66 corresponds to 0.66£ earnings per period). With 4 exceptions, estimated earnings from buying (H, L) instead of (M, L) are also higher, though the number of significant cases, at 5, is much reduced.

The fact that higher compensation is provided for attaining the first best instead of the second-best relative to reaching the second-best instead of the third-best suggests that cognitive effort cost is strictly convex.

7. Conclusion

We find that, when goods are indivisible, markets set prices so that budget allocations become computational sufficiently complex. This makes agents with homogeneous preferences choose different allocations depending on how much cognitive effort they exert. The resulting heterogeneity in demands allows markets to equilibrate. Those who spend more cognitive effort are compensated. This contrasts with traditional Walrasian equilibrium, where existence relies on multiple (indifferent) budget allocation optima, regardless of computational complexity.

We have shown that an equilibrium that compensates for cognitive effort may exist even if the Walrasian equilibrium does not. Conditions under which equilibrium exists need to be re-evaluated. By scattering agents based on cognitive effort, markets restore equilibrium by means of heterogeneity. Heterogeneity in choices (and hence, revealed preferences) is known to be crucial for generic existence of equilibrium under indivisibilities.

We close with a philosophical note. Anxiety has risen recently in modern societies (e.g., Goodwin et al. (2020)). One way markets could cause anxiety is by making budget allocation instances hard: even if substantial effort is applied, the agent is still not sure whether she chooses in the best way. The recent rise in anxiety has coincided with an increased use of market-based solutions to solve societies' problems, such as transportation (airline deregulation, railway privatization), medical care (healthcare marketplaces), education (higher education loan programs), climate change (carbon markets). One wonders whether there is a link.

References

Arora, S. & Barak, B. (2010), *Computational Complexity: A Modern Approach*, Cambridge University Press, Cambridge.

Asparouhova, E. N., Bossaerts, P. & Ledyard, J. O. (2024), Price formation in multiple, simultaneous continuous double auctions, with implications for asset pricing, Technical report.
URL: *http://dx.doi.org/10.2139/ssrn.2774301*

Bossaerts, P., Bowman, E., Fattinger, F., Huang, H., Lee, M., Murawski, C., Suthakar, A., Tang, S. & Yadav, N. (2024), 'Resource allocation, computational complexity, and market design', *Journal of Behavioral and Experimental Finance* **42**, 100906.

Franco, J. P., Bossaerts, P. & Murawski, C. (2024), 'The neural dynamics associated with computational complexity', *PLOS Computational Biology* **20**(9), e1012447.

Franco, J. P., Yadav, N., Bossaerts, P. & Murawski, C. (2021), 'Generic properties of a computational task predict human effort and performance', *Journal of Mathematical Psychology* **104**, 102592.

Gilboa, I., Postlewaite, A. & Schmeidler, D. (2021), 'The complexity of the consumer problem', *Research in Economics* **75**(1), 96–103.

Goodwin, R. D., Weinberger, A. H., Kim, J. H., Wu, M. & Galea, S. (2020), 'Trends in anxiety among adults in the united states, 2008–2018: Rapid increases among young adults', *Journal of psychiatric research* **130**, 441–446.

Meloso, D., Copic, J. & Bossaerts, P. (2009), 'Promoting intellectual discovery: patents versus markets', *Science* **323**(5919), 1335–1339.

Murawski, C. & Bossaerts, P. (2016), 'How humans solve complex problems: The case of the knapsack problem', *Scientific Reports* **6**.

Plott, C. R. & Sunder, S. (1982), 'Efficiency of experimental security markets with insider information: An application of rational-expectations models', *The Journal of Political Economy* pp. 663–698.

Plott, C. & Sunder, S. (1988), 'Rational expectations and the aggregation of diverse information in laboratory security markets', *Econometrica* **56**(5), 1085–1118.

Rasooly, I. (2022), Competitive equilibrium and the double auction, Technical report, Paris School of Economics.

Smith, V. L., Suchanek, G. L. & Williams, A. W. (1988), 'Bubbles, crashes, and endogenous expectations in experimental spot asset markets', *Econometrica: Journal of the Econometric Society* pp. 1119–1151.

Smith, V. L. & Williams, A. W. (1982), 'The effects of rent asymmetries in experimental auction markets', *Journal of Economic Behavior & Organization* **3**(1), 99–116.

Yadav, N., Hsu, A., Franco, J. P., Bossaerts, P. & Murawski, C. (2022), 'How well could and do humans approximate optimality?', *Under Review* .

Yadav, N., Murawski, C., Sardina, S. & Bossaerts, P. (2020*a*), 'Is hardness inherent in computational problems? Performance of human and digital computers on random instances of the 0-1 knapsack problem', *Proceedings of the 24th European Conference on Artificial Intelligence* .

Yadav, N., Murawski, C., Sardina, S. & Bossaerts, P. (2020*b*), Is hardness inherent in computational problems? performance of human and electronic computers on random instances of the 0-1 knapsack problem, *in* 'ECAI 2020', IOS Press, pp. 498–505.