

Sybil Proofness in BRACE: A Mathematical Analysis of Sybils in Competitive Combinatorial Exchanges

Abhimanyu Nag • 17 Nov 2025

ERRATA

The original BRACE mechanism is defined purely for ordinal preferences (\succeq_i) and does not assume the existence of cardinal utility functions. In this blog, we use a Lipschitz cardinal utility representation $V_i(p)$ to quantify the sensitivity of prices and welfare to Sybil perturbations. Formally, we assume

$$|V_i(p) - V_i(p')| \leq L_V \|p - p'\|, \quad \forall i \in I,$$

for some constant L_V . This is not guaranteed by BRACE, and is an additional regularity assumption that allows us to apply first-order perturbation techniques.

The linear price and welfare bounds are valid only under this extra assumption. In markets where ordinal preferences are highly discontinuous or "all-or-nothing," these bounds may fail. Thus, our quantitative Sybil sensitivity results should be interpreted as **conditional** on the existence of a well-behaved Lipschitz cardinalization of agents' ordinal preferences, rather than as a property guaranteed by BRACE itself. My research WIP will deal with it more rigorously.

This exposition is for quick feedback and sanity checks before I conduct more research and finally write the paper. I have made some strong claims without formally writing out proofs and tried to stay true to just the results. Feedback should be sent to contact sources listed on [my website](#).

I recently dove deep into the Sybil Proofness rabbit hole while also reading an amazing paper called "[Competitive Combinatorial Exchanges](#)" by Jantschgi et al.

This gave me the idea to give myself an exercise to try and attack this market "for fun and profit" (as ironically introduced by [Qin et al.](#)) using the intuition I picked up from the literature.

This blog post documents that exercise. I've tried to keep things (in)formal enough to be readable, but also rigorous enough that a researcher familiar with market design or cryptoeconomic security will find the arguments meaningful. A detailed research paper will come out with the same results soon.

If you aren't familiar with this area, the BRACE paper itself does an excellent job explaining the economic motivation behind these ideas. Here, I'll give a compact overview for those who just want a tl;dr of this blog. For our non technical friends,

We study the behavior of BRACE (Budget Relaxed Approximate Competitive Equilibrium) in economies where principals may deploy arbitrarily many Sybil identities. We find that BRACE remains fair and well-behaved even when Sybils appear, but prices and allocations shift proportionally to the number of fake identities. An interesting finding is that BRACE is stable to small attacks but cannot fully prevent strategic manipulation if a single actor controls many identities.

And for our technical peeps,

Modeling Sybils as a perturbation of the type-distribution (principal utility distribution), we establish that BRACE remains well-defined, preserves δ -approximate clearing and maintains ordinal welfare guarantees but induces price deviations of order $\mathcal{O}(\alpha)$, where α is the Sybil infiltration rate. We further show that honest-agent welfare degrades at most linearly in α under Lipschitz preferences, while strategyproofness-in-the-large (SP-L) fails whenever any principal controls a non-vanishing identity mass.

Introduction

The **Budget Relaxed Approximate Competitive Equilibrium (BRACE)** mechanism is a recently developed allocation rule designed for markets where:

- goods are **indivisible**,
- agents only have **ordinal preferences**,
- endowments or rights are **lottery-valued**,
- the system must respect **hard capacity constraints**,
- and equilibrium must exist **even when classical competitive equilibrium fails**.

BRACE introduces a tiny random **budget relaxation**, which transforms non-convex demand (caused by indivisible goods) into convex demand *in expectation*. This small twist is pretty awesome as it restores the existence of equilibrium, while preserving fairness and efficiency guarantees.

Before entering the technical definitions, let's motivate why BRACE is necessary with a simple example.

A Motivating Example

Imagine a platform that needs to allocate two indivisible compute slots, such as two GPU time windows. These slots cannot be split or shared i.e each can only go to a single user, so the platform faces a simple but rigid capacity constraint.

Assume that one unit of Slot A and one unit of Slot B are available in this round. Three users (agents) are competing for these resources and each has straightforward ordinally ranked preferences over the possible outcomes. Agent 1 most prefers to receive Slot A, would accept Slot B if A is unavailable, and considers receiving nothing as the least desirable option. Agent 2 has the exact opposite priorities: they most prefer Slot B, then Slot A, and finally nothing. Agent 3 is less discriminating; they strictly prefer receiving either slot over receiving nothing, but they have no preference between A and B i.e they value both equally.

To make the scenario realistic, assume that the platform uses lottery based priority rights to decide how much “purchasing power” or entitlement each agent brings into the round. Notice that systems like blockchain mempools, priority auctions, and repeated matching platforms sometimes use randomization to allocate rights over time in a fair way. In this example, Agent 1 has a 50% chance of being endowed with Slot A from prior rounds, Agent 2 has a 50% chance of being endowed with Slot B, and Agent 3 arrives with no endowed rights at all. These endowments are probabilistic which creates uncertainty about who can feasibly acquire which slot.

One might hope that a **competitive equilibrium** (defined as *a set of prices for the two slots under which each agent selects their preferred affordable option and markets clear*) could resolve the allocation. However, the classical equilibrium framework breaks down almost immediately.

Because the compute slots are indivisible, agents cannot choose fractional bundles to smooth out their choice. They must either take or leave an entire slot. Preferences are also purely ordinal, meaning agents can rank outcomes but cannot express precise tradeoffs between them. Combined with stochastic

endowments, this makes their budget sets irregular and discontinuous. Moreover, both Agent 1 and Agent 2 have such strong preferences that, even at very high prices, they might still demand the same slot which causes non-convex demand, something **classical equilibrium theory cannot handle**. As a result, there may be no set of prices that simultaneously respects agents' preferences, budgets, and the capacity constraints of the two slots.

Therefore, in markets with indivisible goods, ordinal preferences, and lottery-based endowments, the standard competitive equilibrium tools simply do not work. The mathematical conditions needed for equilibrium such as convex demand, continuous excess demand, and feasible clearing can all fail simultaneously, even in very small markets. It is precisely these kinds of environments that motivate the development of mechanisms like BRACE.

Relevant markets where BRACE is applicable

BRACE is especially relevant to markets with indivisibilities, ordinal input formats, and probabilistic rights. Concrete examples:

1. Mempool / Blockspace Priority Allocation.

- Goods: discrete blockspace slots or inclusion slots in a block.
- Agents: transaction senders with priority credits (e.g., past fee rebates, reputation lotteries).
- Endowments: lottery rights from previous rounds (e.g., rollovers of priority).
- Why BRACE fits: blockspace is indivisible at transaction level; ordinal preferences (urgency ranking) and probabilistic queue-crediting are common.

2. Priority Gas Auctions / Sequencing Rights.

- Goods: positions in a transaction sequence or execution ordering (finite discrete slots).
- Agents: MEV relayers, proposers, sequencers with ordinals on positions and stochastic entitlements.
- Why BRACE fits: discrete ordering, ordinal valuation of positions, randomized entitlements (e.g., proposer rotations).

3. Validator Slot / Committee Assignment Markets (staking ecosystems).

- Goods: validator committee seats or ephemeral slots for processing blocks.
- Agents: validators or delegators with probabilistic staking rewards as endowments.

- Why BRACE fits: seats are indivisible; priority is sometimes allocated via lotteries or stake-based sampling; fairness and security constraints are central.

4. Decentralized Exchange (DEX) Matching with Bundle Constraints.

- Goods: discrete fill-quantities in combinatorial trades or batched auction slots.
- Agents: liquidity takers and makers who hold lottery-like rights (e.g., LP coupons).
- Why BRACE fits: combinatorial bundles and indivisible order blocks with ordinal preferences over execution patterns.

5. Off-chain / Layer-2 Resource Scheduling.

- Goods: finite scheduling slots for compute or storage; agents have priority tokens minted via on-chain events.
- Why BRACE fits: allocation must be fair, respects caps, and often uses ordinal priority signals.

In each of these settings the common features are **discrete resources, ordinal inputs or coarse preference information**, and **probabilistic or history-dependent endowments**. Now that we have touched upon the usefulness of BRACE, let us introduce the mathematical model to describe BRACE rigorously.

BRACE Model

Notations

Following the original paper closely, we introduce some notation:

- Let m be the number of distinct indivisible resource types (goods).
- Let $c = (c_1, \dots, c_m)$ denote the capacity vector: c_j copies of good j are available.
- Let \mathcal{B} denote the finite set of *feasible bundles* (multisets of goods). We write a bundle as $b \in \mathcal{B}$, where b_j is the number of copies of good j in bundle b . We assume $\max_{b \in \mathcal{B}} \|b\|_1 \leq \Delta < \infty$.
- Let I be a finite set of **identities** (agents as reported to the mechanism). The BRACE model will later allow augmentation I^α with Sybils (for now we write I).
- For each identity $i \in I$:
 - $\Psi_i \subseteq \mathcal{B}$ denotes the set of *acceptable bundles* for i .
 - $L(\Psi_i)$ denotes the set of probability distributions (lotteries) supported on Ψ_i . We identify each lottery with a vector $x_i \in \Delta(\Psi_i)$ (probability mass function on Ψ_i).

- $\tilde{e}_i \in L(\Psi_i)$ denotes the *endowment lottery* of i : a (possibly degenerate) random bundle that i initially holds.
- \succeq_i is a weak order (complete and transitive) on $L(\Psi_i)$; agents rank lotteries (ordinal preferences). We assume \succeq_i admits a stochastic-dominant representation when required or else we treat \succeq_i abstractly.

The *economy* is

$$E = (I, \{\Psi_i, \tilde{e}_i, \succeq_i\}_{i \in I}, c).$$

All randomness is over lotteries and budget-relaxation draws. When we write $\mathbb{E}[\cdot]$ it integrates over endowment realizations, budget-relaxations, and any mechanism-level randomization. The BRACE paper's formal definitions also appear in the original paper (Definitions 4–6, Propositions 1–2) although the author's notations differ a bit from my own.

Prices and normalization

- Let prices be a vector $p = (p_1, \dots, p_m) \in \mathbb{R}_+^m$.
- Normalize prices onto the simplex:

$$\Delta^{m-1} := \{p \in \mathbb{R}_+^m : \sum_{j=1}^m p_j = 1\}.$$

For a deterministic bundle $b \in \mathcal{B}$, define $p \cdot b := \sum_{j=1}^m p_j b_j$.

- For a lottery $x \in L(\Psi_i)$, define the *price-value*:

$$v_p(x) := \sum_{b \in \Psi_i} x(b) (p \cdot b).$$

Budgets and δ -relaxation

- Under price p , agent i 's *nominal budget* is

$$B_i(p) := v_p(\tilde{e}_i).$$

Since \tilde{e}_i is a lottery, $B_i(p)$ may be random.

- Fix $\delta > 0$. BRACE introduces a **budget-relaxation** random variable \tilde{b}_i drawn i.i.d. from a distribution $\nu(\delta)$ supported on a small simplex.

- The *relaxed budget* is

$$\tilde{B}_i := B_i(p) + \tilde{b}_i.$$

Taking $\delta \rightarrow 0$ recovers the tight-budget setting where equilibrium generally fails to exist.

Demand correspondence

For fixed prices p and realizations $(\tilde{e}_i, \tilde{b}_i)$:

- The **feasible set** of lotteries is:

$$\mathcal{F}_i(p, \tilde{B}_i) := \{x \in L(\Psi_i) : v_p(x) \leq \tilde{B}_i\}.$$

The **demand correspondence** is:

$$D_i(p, \tilde{e}_i, \tilde{b}_i) := \arg \max_{x \in \mathcal{F}_i(p, \tilde{B}_i)} x \text{ w.r.t. } \succeq_i.$$

Regularity properties from BRACE:

- D_i is nonempty for all p .
- D_i is upper hemicontinuous in p when budgets include \tilde{b}_i .
- Lotteries + budget-relaxation smooth out non-convexities from indivisible goods.

Aggregate expected demand and δ -clearing

- **Aggregate expected demand** for good j :

$$\mathcal{D}_j(p) = \mathbb{E} \left[\sum_{i \in I} \mathbb{E}_{x \sim D_i(p, \tilde{e}_i, \tilde{b}_i)} [x_j] \right].$$

Excess demand:

$$Z(p) = \mathcal{D}(p) - c.$$

δ -approximate clearing condition:

$$Z_j(p) \leq \delta c_j \quad \forall j, \quad p_j > 0 \implies Z_j(p) = \delta c_j.$$

Definition: δ -BRACE

A pair $(p, \{\tilde{x}_i\}_{i \in I})$ is a δ -**BRACE** allocation for economy E if:

1. Each agent selects a most-preferred affordable lottery:

$$\tilde{x}_i \in D_i(p, \tilde{e}_i, \tilde{b}_i).$$

2. Prices $p \in \Delta^{m-1}$ satisfy the δ -clearing conditions.

Notice that BRACE defines an equilibrium **in expectation**: the \tilde{x}_i are random lotteries, and clearing is applied to expected aggregate demand. Refer to the original paper for the best way to get first hand feel for these definitions.

Some Theoretical Facts about BRACE

- **Existence (Proposition 1).**

For any finite economy E and any $\delta > 0$, a δ -BRACE exists.

- **Welfare & fairness (Proposition 2).**

Every δ -BRACE is:

- individually rational,
- ordinally efficient w.r.t. $(1 + \delta)c$,
- endowment-monotone,
- identity-level justified-envy-free.

These are proven in the original paper.

Now that we have developed a formal model for BRACE, let us get down to business.

A Sybil Attack on GPU Slot Market

To see how Sybil identities can distort allocations in a BRACE market, return to the GPU-slot example with two slots, A and B, and three honest agents. Agent 1 prefers A over everything else and has a 50% chance of already owning A. Agent 2 prefers B and has a 50% chance of owning B. Agent 3 is flexible and just wants any slot. In this small, balanced market, BRACE sets prices roughly evenly where Slot A and Slot B end up with similar equilibrium prices because demand is symmetric. Each agent receives a fair lottery allocation based on their preferences and endowments.

Now imagine that Agent 1 is actually a strategic principal who creates five fake Sybil identities. Each Sybil pretends to be a new user who desperately wants Slot A and reports a preference ranking of $A > (\text{everything else}) > \text{nothing}$. None of these Sybils come with any initial endowment. Before the attack, only one out of three agents strongly wanted A and after the attack, six out of eight identities suddenly claim that A is their top choice. From the perspective of the BRACE mechanism, which only sees identities and their reported types, this looks like a dramatic surge in demand for Slot A.

BRACE responds to this new apparent demand by raising the price of Slot A relative to Slot B. Before the attack, Slots A and B may have been roughly equal in price; after adding the five Sybils, Slot A might become, for example, 30–40% more “expensive” in the price simplex. BRACE also reallocates expected demand accordingly: more of the scarce budget slack now flows into the A side of the market, matching the fact that most identities claim to want A first.

Notice that this artificial price increase benefits the attacker. The principal’s real endowment which was a genuine 50% probability of owning Slot A, becomes more valuable when Slot A rises in price. Even if the Sybils themselves receive no actual compute time, their presence helps inflate the market value of the principal’s real asset. When the Sybils cost nothing to create, the principal can use them as “gravity wells” that pull the equilibrium toward one side of the market.

Formal Modelling of Sybil Attack in BRACE

Assumptions (for perturbation method)

Before we start the actual derivation, We lay down some assumptions. To derive stability bounds, we typically assume the following:

- **Finite bundle system:** \mathcal{B} finite, $\Delta < \infty$.
- **Bounded budgets:** $B_i(p)$ uniformly bounded; \tilde{b}_i supported in an $O(\delta)$ simplex.
- **Lipschitz aggregate demand:** $p \mapsto \mathcal{D}(p)$ is Lipschitz in the population type measure.
- **Lipschitz indirect utilities:** A cardinal utility representation consistent with \succeq_i yields Lipschitz $V_i(p)$.

These assumptions appear implicitly in BRACE’s continuity and fixed-point arguments.

Principals and Sybil Identities

Definition 1 (Principals)

Let P be the set of **principals** i.e real economic actors. Each identity $i \in I$ belongs to exactly one principal through a mapping

$$\pi : I \rightarrow P.$$

Definition 2 (Sybil Identity)

An identity $i \in I$ is a **Sybil** if

$$|\pi^{-1}(\pi(i))| > 1.$$

Let

$$S = \{i \in I : |\pi^{-1}(\pi(i))| > 1\}, \quad H = I \setminus S,$$

and define the **Sybil infiltration rate**

$$\alpha = \frac{|S|}{|I|}.$$

Definition 3 (Sybil-Augmented Economy)

Given the original economy E , the Sybil-augmented economy is

$$E^\alpha = (I^\alpha, \tilde{E}^\alpha, \succeq^\alpha, \Psi^\alpha, c), \quad I^\alpha = H \cup S.$$

A principal selects $(\Psi_i, \tilde{e}_i, \succeq_i)$ for each of its identities.

BRACE in the Sybil Economy

A δ -BRACE in E^α consists of:

- price vector $p^\alpha \in \Delta$
- budget relaxations $\tilde{b}_i \sim \nu(\delta)$
- random allocations \tilde{x}_i^α

Recall in context:

Expected excess demand:

$$Z_j^\alpha(p) = \mathbb{E} \left[\sum_{i \in I^\alpha} x_{i,j}^\alpha(p) \right] - c_j.$$

BRACE clearing requires

$$Z_j^\alpha(p) \leq \delta c_j, \quad p_j > 0 \Rightarrow Z_j^\alpha(p) = \delta c_j.$$

Theoretical Results

Existence Under Sybils

For every $\delta > 0$, the Sybil-augmented economy E^α admits a δ -BRACE.

Proof.

From the paper, the BRACE existence theorem holds for all finite economies. Sybil augmentation preserves finiteness. \square

ON THE SIDE: What if finiteness condition was broken? It will affect our proofs adversely. Napkin math gives :

$$\left(|I| = \infty \wedge \exists p \in \Delta^{m-1} : \sum_{i \in I} \mathbb{E}[\|D_i(p, \tilde{e}_i, \tilde{b}_i)\|] = \infty \right) \vee (|\mathcal{B}| = \infty)$$

$$\implies \neg \exists (p, \tilde{X}) \text{ that is a } \delta\text{-BRACE for } E.$$

So If a principal creates infinitely many Sybils with identical types, then no price vector $p \in \Delta^{m-1}$ can satisfy the δ -clearing conditions. *This unsurprisingly implies that the first defense is to make sybils costly to create.*

Welfare Validity Under Sybils

Every δ -BRACE in E^α is individually rational and ordinally efficient w.r.t. $(1 + \delta)c$.

Proof. Direct from Proposition 2 of the BRACE paper. \square

Price Perturbation Bounds

Let μ^0 be the empirical type distribution in the honest economy, μ^α the Sybil-perturbed distribution. Let p^0 and p^α denote the corresponding δ -BRACE prices.

We strengthen the previously stated result.

Assumptions

1. \mathcal{B} finite with $\max_b \|b\|_1 \leq \Delta$.
2. **Lipschitz excess demand:** for some $L_Z > 0$,

$$\|Z(p, \mu) - Z(p, \mu')\| \leq L_Z W_1(\mu, \mu')$$

for all p in the simplex.

3. **Jacobian nondegeneracy:** the differential of excess demand at equilibrium satisfies

$$\|(D_p Z(p^0, \mu^0))^{-1}\| \leq \Gamma.$$

Sybil–Price Stability Bound

There exist constants $K = \Gamma L_Z$ and $\bar{\alpha} > 0$ such that for all $\alpha \leq \bar{\alpha}$,

$$\|p^\alpha - p^0\| \leq K\alpha.$$

Proof

We supply a proof using perturbation of roots (took a while to deal with).

Let

$$F(p, \mu) := Z(p, \mu) - \delta c.$$

The BRACE clearing condition requires $F(p, \mu) = 0$.

By assumption,

$$F(p^\alpha, \mu^\alpha) = 0, \quad F(p^0, \mu^0) = 0.$$

Consider a first-order expansion around (p^0, μ^0) :

$$0 = F(p^\alpha, \mu^\alpha) = F(p^0, \mu^0) + D_p F(p^0, \mu^0)(p^\alpha - p^0) + D_\mu F(p^0, \mu^0)(\mu^\alpha - \mu^0) + R,$$

where $\|R\| = o(\|p^\alpha - p^0\| + W_1(\mu^\alpha - \mu^0))$.

Since $F(p^0, \mu^0) = 0$, rearrange:

$$D_p F(p^0, \mu^0)(p^\alpha - p^0) = -D_\mu F(p^0, \mu^0)(\mu^\alpha - \mu^0) - R.$$

Invert using assumption 3:

$$p^\alpha - p^0 = -(D_p F(p^0, \mu^0))^{-1} D_\mu F(p^0, \mu^0)(\mu^\alpha - \mu^0) - (D_p F)^{-1} R.$$

Taking norms:

$$\|p^\alpha - p^0\| \leq \Gamma L_Z W_1(\mu^\alpha, \mu^0) + o(\alpha).$$

Since $\mu^\alpha = (1 - \alpha)\mu^0 + \alpha\mu^S$,

$$W_1(\mu^\alpha, \mu^0) \leq \alpha W_1(\mu^S, \mu^0) \leq \alpha.$$

Thus,

$$\|p^\alpha - p^0\| \leq K\alpha$$

for $K = \Gamma L_Z$, for sufficiently small α . \square

Honest Welfare Bounds

Let the indirect utility of an honest identity be $V_i(p)$.

Yet Another Assumption

Each V_i is L_V -Lipschitz:

$$|V_i(p) - V_i(p')| \leq L_V \|p - p'\|.$$

Let $W_H(E)$ be the total welfare of honest agents.

Honest Welfare Loss Bound

For sufficiently small α ,

$$W_H(E^\alpha) \geq W_H(E) - |H|L_V K\alpha.$$

Proof

Using Lipschitz utilities and Price Stability Bound,

$$|V_i(p^\alpha) - V_i(p^0)| \leq L_V \|p^\alpha - p^0\| \leq L_V K\alpha.$$

Summing over honest identities:

$$W_H(E^\alpha) = \sum_{i \in H} V_i(p^\alpha) \geq \sum_{i \in H} V_i(p^0) - |H|L_V K\alpha = W_H(E) - |H|L_V K\alpha.$$

\square

Therefore, in the presence of a small perturbation α , the honest agents' total welfare will only decrease by a small amount, proportional to the size of the perturbation constants.

Incentives and Failure of Strategyproofness in the Large (SP-L)

Now to deal with the strongest result. Define $C_p = \pi^{-1}(p)$, the set of identities belonging to principal p .

Necessary and Sufficient Conditions for SP-L

BRACE is strategyproof-in-the-large **iff**

$$\max_{p \in P} \frac{|C_p|}{|I|} \rightarrow 0.$$

Proof

(\rightarrow necessity)

Suppose some principal p has $\gamma := \limsup |C_p|/|I| > 0$.

Manipulating the types of γn identities changes the empirical distribution by

$$W_1(\mu^\alpha, \mu^0) \geq \gamma,$$

so by Stability Bound:

$$\|p^\alpha - p^0\| \geq K\gamma.$$

Thus the principal can change prices by a non-vanishing amount thereby violating negligible influence.

Hence SP-L fails.

(\leftarrow sufficiency)

If

$$\max_p \frac{|C_p|}{n} \rightarrow 0,$$

then every principal affects the empirical distribution by $o(1)$, implying via Stability Bound that

$$\|p^\alpha - p^0\| = o(1),$$

which is exactly the negligible-influence condition under which SP-L holds (BRACE Theorem 4).

□

Therefore, BRACE is strategyproof-in-the-large exactly when no single principal can control a significant fraction of identities. Large numbers dilute influence and ensures negligible impact and no incentive to misreport.

Principal Fairness

Identity-level Justified Envy-Freeness (JEF) does not imply principal-level fairness.

Principal-Level JEF Failure

There exist economies E^α and principals p, q such that:

- every identity satisfies identity-level JEF
- but at principal level:

$$x_p \prec_{\text{JEF}} x_q.$$

Proof

Construct a principal p with many Sybils $\{i_1, \dots, i_k\}$ whose endowments are negatively correlated, while another principal q holds a single identity.

BRACE allocations are JEF **for each identity**, but principal p can aggregate $\tilde{x}_{i_1}, \dots, \tilde{x}_{i_k}$ into a diversified portfolio with lower variance.

Since BRACE fairness is defined on identities, the aggregation step

$$x_p := \sum_{i \in \pi^{-1}(p)} \tilde{x}_i$$

is **outside** the mechanism's fairness guarantees.

Thus x_p may dominate x_q in every JEF comparison. \square

Incentive Compatible Sybil-Proofness Bounds

Here give **explicit bounds** on how much a principal can gain by injecting Sybils.

Let the Sybil-augmented distribution be

$$\mu^\alpha = (1 - \alpha)\mu^H + \alpha\mu^S.$$

Let p^α be the BRACE price.

Define the **principal's utility**:

$$U_p(\alpha) := \sum_{i \in C_p} V_i(p^\alpha).$$

Marginal Sybil Incentive Bound

For sufficiently small α ,

$$\frac{d}{d\alpha} U_p(\alpha) \leq |C_p| L_V K.$$

Proof

By chain rule,

$$\frac{d}{d\alpha} U_p(\alpha) = \sum_{i \in C_p} \langle \nabla V_i(p^\alpha), \frac{dp^\alpha}{d\alpha} \rangle.$$

Use Lipschitz gradient bound:

$$\|\nabla V_i(p)\| \leq L_V.$$

Use Stability Bound:

$$\left\| \frac{dp^\alpha}{d\alpha} \right\| \leq K.$$

Thus

$$\frac{d}{d\alpha} U_p(\alpha) \leq |C_p| L_V K.$$

□

Hence a corollary

Sybil Proofness at Rate α

If a principal controls k Sybils, the **maximum gain** from Sybil creation is:

$$U_p(\alpha) - U_p(0) \leq k L_V K \alpha.$$

Therefore, the additional utility a principal can obtain by creating k Sybils is at most $kL_V\alpha$, growing **linearly** with the number of Sybils.

A Stronger Impossibility Result

Continuing from the above, we provide a new result: unless BRACE is modified, *no competitive equilibrium with indivisible goods can be Sybil-proof beyond linear order.*

No Better Than $O(\alpha)$ Sybil Resistivity

For any δ -BRACE mechanism and any $\epsilon > 0$, there exist economies where a principal can achieve welfare gains $\Omega(\alpha)$, and **no mechanism satisfying BRACE axioms** can reduce this below $o(\alpha)$.

Proof (Sketch)

The tightness of Stability Bound follows from a first-order perturbation lower bound:

$$\|p^\alpha - p^0\| \geq c\alpha$$

for some $c > 0$ whenever the principal's Sybil strategy induces type variation of order α .

Utilities then shift by $\Omega(\alpha)$ under Lipschitzness.

Thus $O(\alpha)$ is both upper **and** lower tight.

□

An Ambitious Attempt at Distributional Sybil-Proofness

I also played around and thought to try and derive a **measure-theoretic** condition for full Sybil-proofness.

Let the type distribution be absolutely continuous w.r.t. a reference measure λ with density f .

Radon–Nikodym Sybil-Proofness Criterion

BRACE is Sybil-proof **iff** every principal’s manipulation satisfies:

$$\frac{d\mu^\alpha}{d\lambda} = \frac{d\mu^0}{d\lambda} \quad \text{a.e.}$$

Proof

(\rightarrow) If the equality holds, the empirical distribution does not change, implying $p^\alpha = p^0$.

(\leftarrow) If the densities differ on a positive-measure set, then $W_1(\mu^\alpha, \mu^0) > 0$, implying by Stability Bound that $p^\alpha \neq p^0$, hence a profitable deviation exists.

□

This result shows that BRACE’s Sybil-proofness is equivalent to the empirical distribution of principals remaining exactly unchanged under manipulation. Mathematically, it’s elegant because it reduces a strategic, combinatorial attack resistance to a precise condition on Radon–Nikodym derivatives. I am personally obsessed with this since it links mechanism design directly to measure theory.

Conclusion

Now as usual I will let ChatGPT give you all a nice send off:

In this article, we extended the BRACE framework to economies in which principals may replicate themselves through Sybil identities. We began by formally distinguishing between *identities* and *principals*, and by allowing each principal to choose arbitrary preference and endowment types for the identities it controls. Within this augmented model, we proved that the core structural guarantees of BRACE—existence, individual rationality, ordinal efficiency with respect to $(1 + \delta)c$, and identity-level justified-envy-freeness—continue to hold. This follows from the fact that BRACE requires only the finiteness of the population and the regularity induced by lotteries and budget-relaxations; Sybil augmentation preserves these properties.

The main contribution of our analysis lies in quantifying the sensitivity of BRACE prices and allocations to Sybil infiltration. By studying the BRACE fixed-point map under perturbations of the empirical type distribution, we established linear price-continuity bounds: if Sybils constitute an α -fraction of the population, the induced equilibrium price vector p^α satisfies

$$\|p^\alpha - p^0\| = O(\alpha).$$

Under mild Lipschitz assumptions on indirect utilities, this yields correspondingly linear distortions in honest-agent welfare. These results highlight a key dichotomy: BRACE is *stable* under small identity perturbations, but not *robust* to adversarial manipulation when a principal controls a non-negligible portion of the population.

Our analysis further revealed two structural limitations. First, BRACE is not strategyproof-in-the-large when any principal's share of the population is bounded away from zero; such principals can manipulate the empirical type distribution sufficiently to affect the BRACE fixed point. Second, BRACE's fairness guarantees apply at the identity level, not the principal level. Once individual allocations are aggregated back to principals, justified-envy-freeness may fail, providing avenues for unfair advantage through identity replication.

Would love thoughts, comments and suggestions.

I have been [Abhimanyu Nag](#)