Quadratic funding under constrained budget is suboptimal

Old School Mathematicians
old@email

August 8, 2025

Abstract

We analyze the efficiency of Quadratic Funding (QF) under individual budget constraints, extending the framework of Weyl et al. (2019). QF's optimality requires participants to equalize marginal utilities across projects—an assumption that fails when budgets are binding, as shown through utility function analysis and a proof by contradiction. Using 5.2 million donations from 23 Gitcoin Grants rounds and multiple Octant Epochs, we document a shift from early donors supporting multiple projects with larger amounts to recent cohorts making small, single-project contributions. These patterns produce peaked preference distributions, a power law in projects supported, and an exponential decay in donation amounts. Our results indicate that, under realistic budget limits and concentrated preferences, QF may deviate from its theoretical efficiency, underscoring the need for mechanism designs that address behavioral and informational constraints.

Keywords Quadratic funding · Budget constraints · Utility optimization

1 Quadratic Funding Definition

QF is defined by Weyl et al. (2019)

2 Logic Flow

An individual i with a budget constraint K_i is trying to donate to his preferred projects from the set P of all projects in a given fundraising round. In Weyl 2019, the optimal amount to donate c_i^p for project $p \in P$ is determined by maximizing his utility

$$U_i^p(F_p) - c_i^p \tag{1}$$

where $F_p = \sum_j c_j^p$ is the total funding for project p from all individuals and U_i^p is the utility that individual i derives from the project p. Weyl et al maximizes this expression by setting the first derivative to 0

$$\begin{aligned} &\frac{dU_i^p}{dc_i^p} - 1 = 0 \\ \Leftrightarrow &\frac{dU_i^p}{dF^p}\frac{dF^p}{dc_i^p} = 1 \end{aligned} \tag{2}$$

In other words, the individual i puts in more money until his utility for the project p is saturated, thereby receiving less than a unit dollar's worth of utility for the next unit dollar put in. A critical assumption implicit in this work by Weyl et al is that K_i is large enough to saturate all projects considered:

$$\sum_{p \in P_i} c_i^p \leq K_i \tag{3}$$

where $P_i \subset P$ is the set of projects that the individual i invests in.

This unstated assumption in Weyl 2019 is critical to the proof of optimality therein. In practice, a single individual do not have budgets to saturate every project considered in a fundraising round. Therefore, quadratic funding optimality condition is not proven under a budget constraint.

We make a different claim.

Lemma 1 An individual i has his total utility $\sum_p U_i^p$ maximized under the following necessary and sufficient conditions:

$$\begin{gathered} \frac{dU_i^a}{dc_i^a} = \frac{dU_i^b}{dc_i^b} \\ \frac{dU_i^l}{dc_i^l} \leq \frac{dU_i^a}{dc_i^a} \\ \forall a, b \in P_i \subset P \qquad \forall l \in P \setminus P_i \end{gathered} \tag{4}$$

In other words, the projects invested must have the same marginal utility; the project uninvested must have lower marginal utility than any of the invested projects.

Proof We prove by contradiction. For the first equality (utility parity), we assume that the equality is not satisfied but the total utility is still maximized. In this case, we can extract one dollar from project $a \in P_i$ that has lower marginal utility and invest the extracted dollar into a project $b \in P_i$ with higher marginal utility. As the decrease in U_i^a is smaller in magnitude than the increase in U_i^b, the total utility for individual i has increased. This contradicts the initial assumption that total utility is maximized.

For the second inequality, we assume that inequality is false. Then, there is an uninvested project offering higher marginal utility, thereby warranting the marginal dollar investment by divesting from another invested project. This is a contradiction, thus proving the necessary case.

We can use a similar argument and the concavity of utility functions to argue the sufficient case. □

We note that the optimality condition presented above allows for single-donation individuals more gracefully than the original QF framework proposed in Weyl 2019. In the original paper, a single-donation is possible when

$$\begin{gathered} \frac{dU_i^a}{dc_i^a} = 1 \\ \frac{dU_i^l}{dc_i^l} < 1 \\ \{a\} = P_i \subset P \qquad \forall l \in P \setminus P_i \end{gathered} \tag{5}$$

whereas with our optimality conditions,

$$\frac{dU_i^l}{dc_i^l} < \frac{dU_i^a}{dc_i^a}$$
$$\{a\} = P_i \subset P \qquad \forall l \in P \setminus P_i \tag{6}$$

which is much more general and models the reality better. We show the prevalence of single-donation individuals in the empirical section below.

3 Empirical Study

To empirically characterize donor utility functions, we analyze the distribution of project support patterns across 5.2 million donations spanning 23 Gitcoin Grants rounds and multiple Octant Epochs. Our analysis reveals several key trends in donor behavior over time.

3.1 Trends

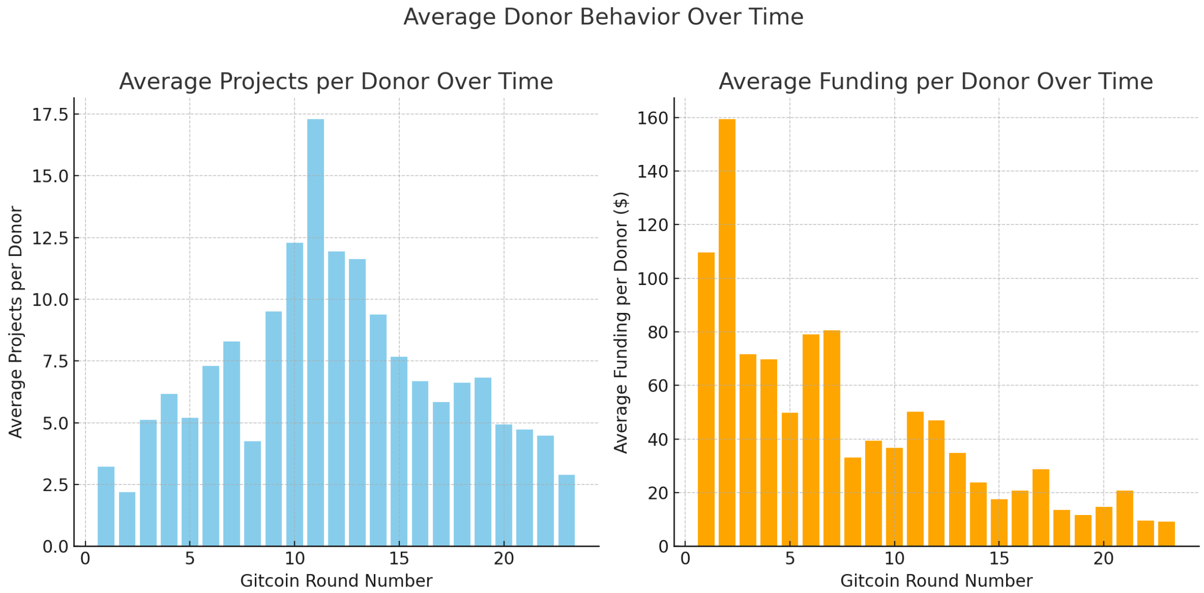

Figure 1: Gitcoin avg funding and supported number of projects.

3.1.1 Evolving Donor Behavior and Mechanism Design

Since Gitcoin Grant Round 11 (GG11), we observe a systematic decline in both average funding amounts and average project support per donor. The distribution of projects supported per donor exhibits power law characteristics:

$$f(x) = Cx^{\{-\alpha\}} \tag{7}$$

while individual donor funding amounts follow an exponential decay distribution:

$$f(x) = \lambda e^{\{-\lambda x\}} \tag{8}$$

3.1.2 Shifting Participant Demographics

We attribute these distributional changes to evolving marketing strategies that have successfully expanded the donor base but fundamentally altered engagement patterns. Recent cohorts demonstrate markedly different behavior: new participants typically contribute small amounts ($1-10) to single projects, contrasting sharply with earlier donors who supported multiple projects with larger individual contributions.

The influence of outreach strategies and information asymmetries becomes particularly evident in recent rounds. Over 58% of donors in GG23 supported only a single project, yet this majority appears unaware that Gitcoin implemented a new matching algorithm (COCM) starting in GG22. This algorithmic transition fundamentally changed how contributions generate matching funds, meaning that single-project donations—while still valuable—no longer optimize matching efficiency under the updated mechanism.

These findings reveal a critical gap: donor education has not kept pace with user acquisition efforts, potentially undermining the intended efficiency gains of quadratic funding systems.

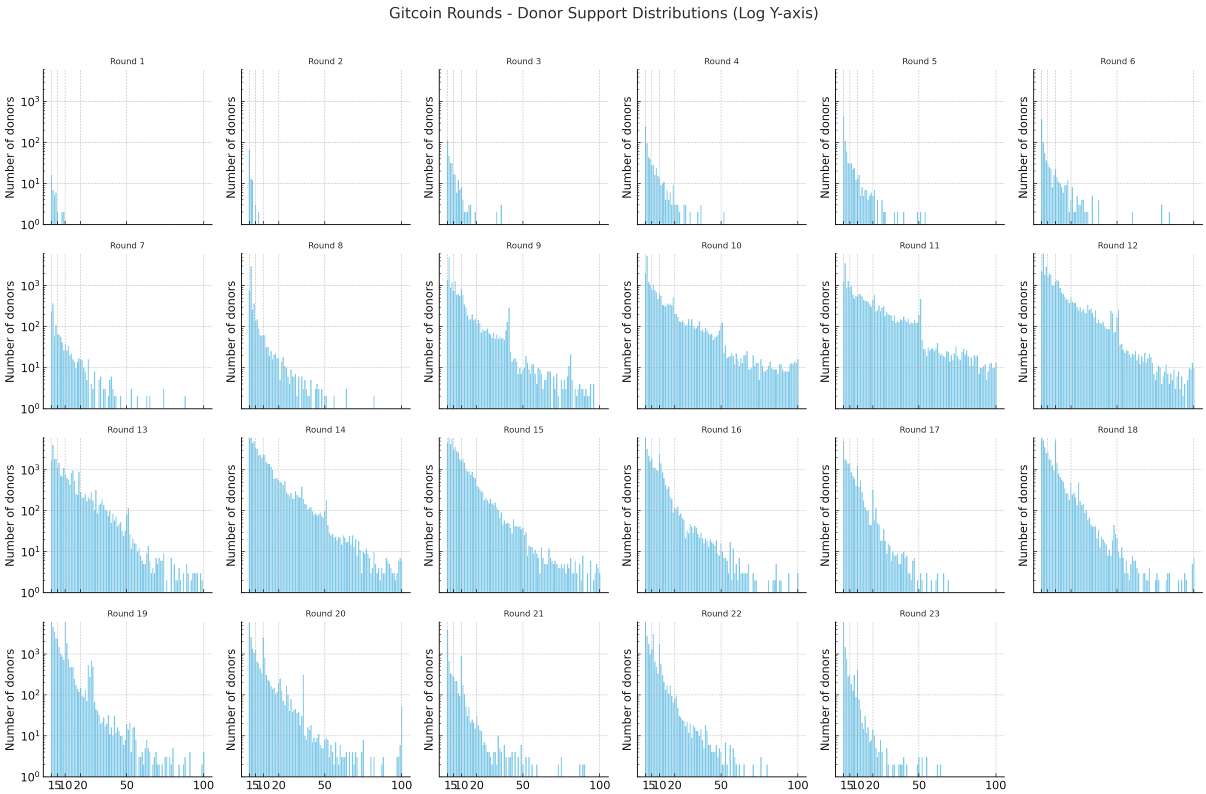

Figure 2: Gitcoin donor's distribution over a period of 23 rounds.

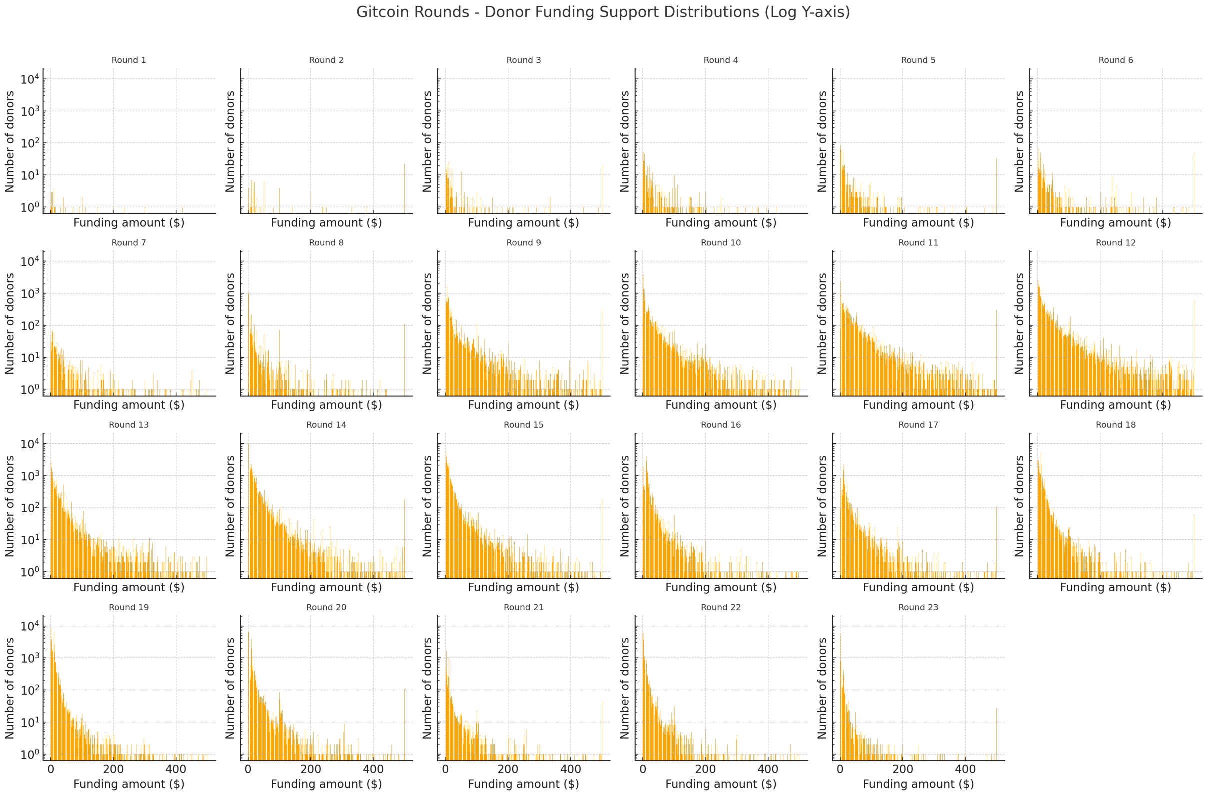

Figure 3: Gitcoin donation distribution over a period of 23 rounds.

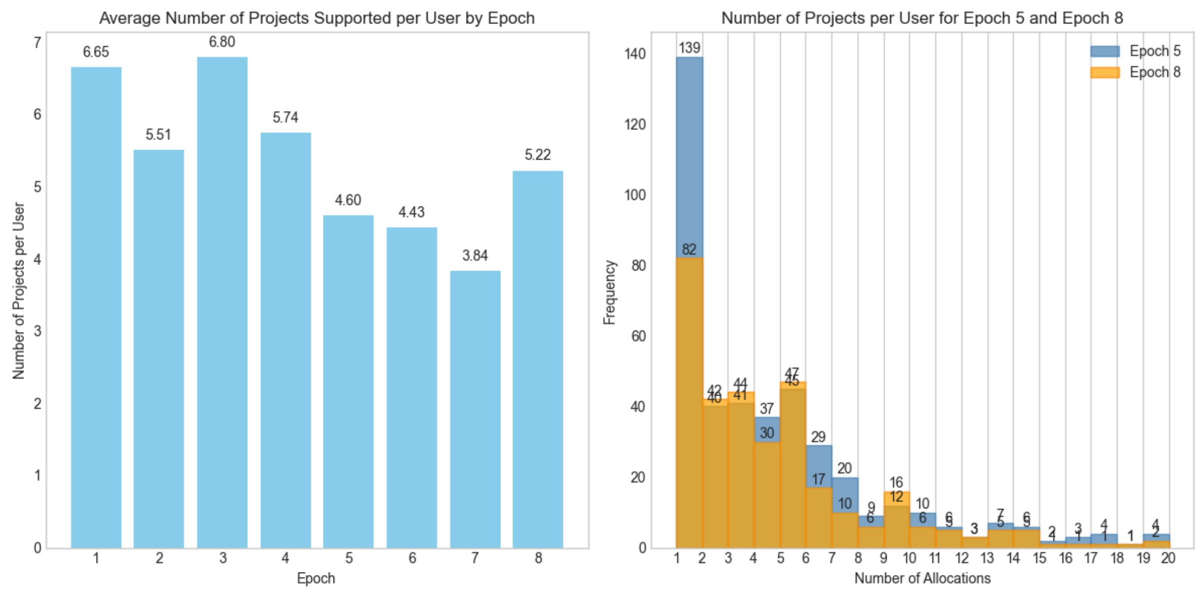

Figure 4: Average number of projects supported by donors during Octant's Epochs.

3.2 On Utility Assumptions

Our empirical findings challenge fundamental assumptions about utility structures in Quadratic Funding (QF) mechanisms. The observed concentration of donations on single projects per funding round suggests highly peaked preference distributions rather than the distributed utility profiles typically assumed in QF theory.

This pattern implies that most donors assign negligible utility to the vast majority of available projects, allocating their entire budget to one preferred option. While consistent with strong preference ordering, this behavior contradicts QF's theoretical foundation, which requires participants to have positive but varying utilities across multiple projects to achieve optimal public goods allocation.

The prevalence of single-project funding raises fundamental questions: do these patterns reflect genuine utility functions, or are they artifacts of information constraints, cognitive limitations, or misaligned incentives? If users truly derive utility from only one project per round, QF mechanisms may not be operating within their intended theoretical regime, potentially compromising their efficiency properties for public goods provision.

3.3 Exponential Decay or Power Law?

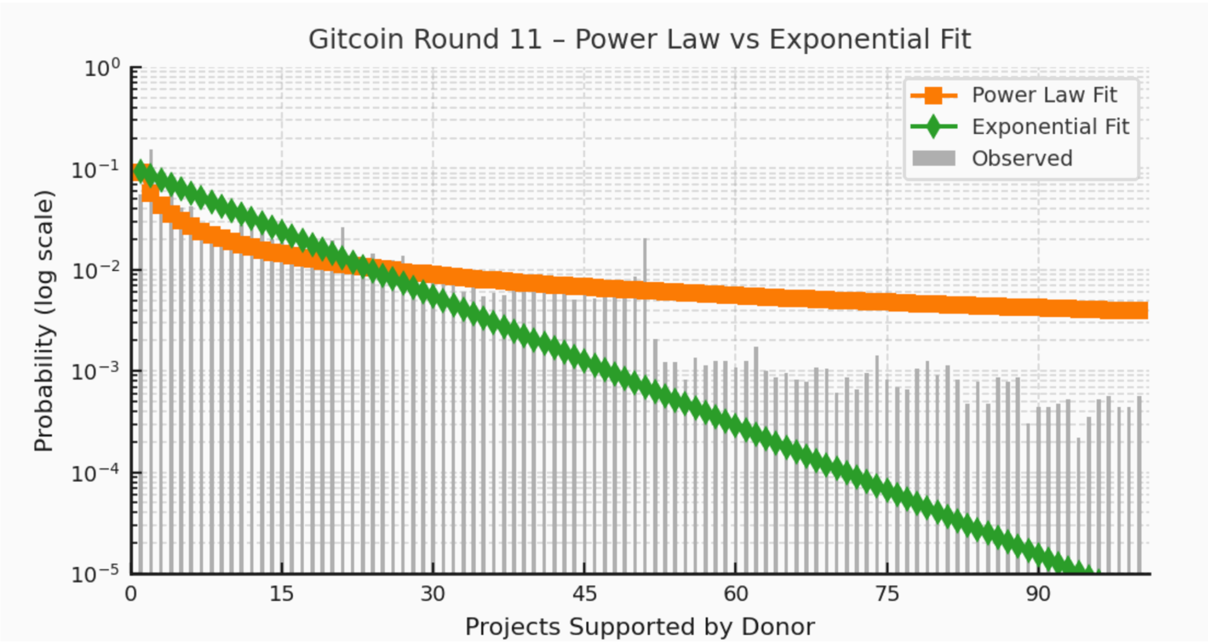

Figure 5: Fitting exponential decay (an extreme case of the poisson distribution) and power law.

3.3.1 Power Law Distribution

Power laws are remarkably common across natural and social systems. They appear in city population distributions (Zipf's law), income and wealth inequality (Pareto distribution), earthquake magnitudes (Gutenberg-Richter law), word frequency in languages, internet network connectivity, and even the distribution of crater sizes on planetary surfaces. This ubiquity stems from underlying processes involving preferential attachment, multiplicative growth, or optimization under constraints. The scale-free nature of power laws—meaning they look the same at different magnifications—makes them particularly important for understanding systems without characteristic scales, from biological networks to financial markets.

3.3.2 Exponential Decay

While the distribution of projects supported per donor cannot be modelled through a Poisson distribution, the funding distribution does fit an exponential decay curve (which connects to the Poisson distribution for small λ values).

The λ has the useful interpretation of likelihood of event happening. In donation amount distribution that means that **λ represents the rate parameter controlling how quickly donation probabilities decay with increasing amounts**. A small λ indicates that large donations are relatively more likely (slower decay), while a large λ suggests that most donations cluster around smaller amounts with rapidly declining probabilities for larger gifts.

This parameter effectively captures donor behavior patterns: λ reflects both the typical scale of donations and the community's propensity for larger contributions. When λ is small, we observe a more gradual decline in donation frequency across amount ranges, suggesting either a wealthier donor base or stronger motivation for substantial giving. Conversely, large λ values indicate a steep drop-off after small donations, characteristic of grassroots fundraising or communities with limited resources.

The exponential decay pattern thus could provide a signal for understanding and predicting donation distributions, providing insights in a funding community's giving behavior.