# What Kinds of Incentives Encourage Participation in Democracy? Evidence from a Massive Online Governance Experiment*

Andrew B. Hall†
Eliza R. Oak‡

November 14, 2023

**Abstract**

How can we democratically govern the AI, social media, and online platforms of the future? Today, low participation is a major barrier to community governance online. We leverage a digital quasi-experiment that allows us to study the links between incentives and political participation at a scale and granularity not previously possible. We focus on a web3 startup called Optimism that distributed digital tokens worth roughly $28 million USD to more than 300,000 active participants in its governance system as part of a sequence of rewards meant to encourage long-run engagement. Studying 1.2 million unique wallet addresses, we find that the reward scheme induces new users to participate in Optimism's governance, has larger effects for smaller token holders, and leads voters to spread their votes across delegates more widely. Together, the results suggest that reward schemes that give people a durable stake in the community and promise a sequence of future rewards are able to broaden participation in online democracy noticeably, at least in the short run under the proper conditions.

*Authors are listed in alphabetical order and contributed equally. The authors thank Daren Matsuoka for invaluable help with data. For helpful discussion, the authors thank Jennifer Gandhi, Greg Huber, Connor Huff, Neil Malhotra, Ethan Bueno de Mesquita, Pablo Montagnes, Tim Roughgarden and participants in the a16z crypto research seminar.

†Davies Family Professor of Political Economy, Graduate School of Business, Stanford University. andyhall@stanford.edu. Hall receives consulting income from a16z crypto and Meta Platforms, Inc.

‡Ph.D. Student, Yale University. eliza.oak@yale.edu. Oak receives consulting income from a16z crypto.

# 1 Introduction

As online platforms play a larger and larger role in society, efforts to democratize them are gaining steam. In response to concerns they hold too much power over consequential decisions that affect society, Meta, OpenAI, and Twitter (now, X) have all recently announced efforts to make a small number of important decisions with varying degrees of democratic input.[1] In parallel, online communities in web3— internet platforms built on top of blockchains, which we will explain below—are experimenting with building democratic governance systems into their platform architecture.

Democratizing online platforms is difficult. Users of online systems expect convenience and are generally uninterested in participating in governing the platforms that they use. Rates of voting in online communities in the web3 space are generally quite low (Barbereau et al. 2022; Fritsch and Wattenhofer 2022; Messias and Loiseau 2023). This reflects a long history in online governance; for example, when Facebook attempted to hold a global referendum of its user base in 2009, it was forced to cancel the effort after less than 1% of all users voted.[2] Faced with the well-known problem of encouraging participation in democracy,[3] online platforms are now experimenting with a variety of ways to incentivize participation, including directly rewarding people for participating. But how can these incentives be designed to effectively encourage participation? And what can they teach us about the nature of political participation more broadly?

An important literature in political science studies the causal effects of paying people to vote, generally finding positive though often modest effects (La Raja and Schaffner 2022; Panagopoulos 2013; Shineman 2018).[4] In online governance, however, there is potentially a broader design space for incentives that tap into other ways of motivating voters but that have not been tried yet in the

[1] For Meta's announcement, see `https://about.fb.com/news/2022/11/improving-peoples-experiences-through-community-forums/`. For OpenAI, see `https://openai.com/blog/democratic-inputs-to-ai`. And for Twitter, see `https://help.twitter.com/en/using-twitter/community-notes` and `https://apnews.com/article/elon-musk-biden-twitter-inc-technology-congress-d88e3de4b3cc095926dc133f53dc3320`.

[2] `https://www.latimes.com/archives/blogs/technology-blog/story/2009-04-23/facebook-governance-vote-is-a-homework-assignment-no-one-did`

[3] A large academic literature has sought to document when and why individuals engage civically. Scholars have investigated, for instance, the role of individual-level social status, education, and income on political efficacy (Highton and Wolfinger 2001; Verba and Nie 1972), which institutional mechanisms (i.e. electoral systems or legislative structures) maximize turnout (Blais and Dobrzynska 1998; Jackman 1987; Lijphart 1997), the effects of inducing social benefits or peer pressure on participation (Bryan et al. 2011; Gerber, Green, and Larimer 2008), or the impact of various Get-Out-The-Vote efforts to directly contact voters (Green and Gerber 2015).

[4] A related literature finds large effects of compulsory voting on turnout (Fowler 2013). Since compulsory voting is typically enforced via relatively small fines, it could suggest a relatively large effect of small payments on turnout, but evidence suggests norm-setting and non-monetary incentives probably drive the effect (Gonzales, León-Ciliotta, and Martínez 2022).

physical world and therefore haven't been studied empirically. For example, what if the rewards voters receive for participating are not purely monetary but instead give them a stake in the future of the community? And what if voters are rewarded over time for sustained participation, not just in a single shot?

In this paper, we leverage a novel, large-scale natural experiment in the digital world to study these questions.[5] We study a startup called Optimism, which allocated rewards to several hundred thousand community members based on their participation in collective governance. Instead of giving users money in exchange for their participation, Optimism allocates rewards in the form of digital "OP tokens" that give users a potential stake in the future of the project. These OP tokens can be used within the Optimism system, convey irrevocable voting rights, and can also be sold on the secondary market. This is an especially interesting intervention to test, as a result, because of existing evidence that people with more at stake economically are more likely to participate in elections (Marble and Nall 2021; Yoder 2020) and to engage in political cooperation (Jha and Shayo 2019).

In addition, instead of providing a single-shot reward for participating, Optimism has promised an ongoing sequence of rewards, announcing these "airdrops"—so called because the rewards are automatically transferred to each user's digital wallet—on a regular basis, keeping secret precisely what governance actions will be rewarded and by what formula until the moment of each announcement (at which point the rewards are immediately transferred to users). As a result, each airdrop announcement not only rewards those who have participated but potentially conveys information to the whole community about what actions will be rewarded in future airdrops.

Finally, because Optimism uses blockchain technology, as we will explain below, all of the data concerning users and their rewards is public—including who received the rewards, how much they received, whether they sold or kept them, whether they used them to vote, and if so, how they voted. As such, the experiment provides a unique chance to evaluate a large-scale intervention on political participation with an unprecedented wealth of individual-level data that would not be possible in regular democratic elections.

[5]This builds on recent work that uses experiments on online platforms to study impacts on political behavior in the physical world (e.g., Allcott et al. 2020; Bond et al. 2012; Guess et al. 2023).

Using a regression discontinuity design, we find that the announcement of Optimism's governance-focused airdrop caused a very large increase in subsequent governance participation, as users inferred that participation would be further rewarded in the future. Users who directly received rewards and therefore had a larger stake in Optimism responded by participating more, on average, even though they had the option to sell the rewards on the open market. Effects were larger for smaller tokenholders, and users were successfully encouraged to delegate their voting power to a broader range of potential governance representatives than in the past. Overall, the intervention appears to have broadened democratic participation within Optimism's system, for the time being at least.

While Optimism is a unique case, at the end of the paper we discuss what our findings suggest about how to design effective online governance for AI, social media, and other kinds of technology platforms, and about the nature of political participation more generally. For online communities, our findings suggest that designing reward schemes that use internal tokens or other kinds of point systems and allocate rewards over a long period of time may be a promising way to encourage participation, if users find the rewards valuable and believe the promises regarding future rewards. Beyond online governance, the results suggest that policies intending to reward participation might be more effective if rewards are not solely monetary but rather give recipients a long-run stake in the community.

## 2 Optimism, Airdrops, and web3 as a Governance Laboratory

The web3 context is a novel laboratory for studying governance, and ours is to our knowledge the first paper in political science to leverage an experiment of this kind. As such, it is important to explain how web3 governance works, the context of the specific experiment we study, and why it is a valuable and generalizable case to analyze.

### 2.1 Background on web3 Governance

The term "web3" is used by supporters to describe a third phase of the internet, one built on decentralized blockchains[6] and where platform ownership is distributed among builders, operators,

[6]The key takeaway about blockchains for this paper about online governance is that they are a new technology meant to offer a more decentralized version of the internet. The desire for decentralization leads to interesting

and users via digital tokens. Projects in web3—alternatively described as companies, platforms, protocols, or startups—typically center around a public repository of open-source code that enables a particular service, which could include things like a social media platform, a financial service provider, a crowdsourced investment group, an art project, or many other applications.

Voting is essential to web3. Projects in web3 are typically governed through a body called a "Decentralized Autonomous Organization" or DAO, where token holders collectively cast votes. Focusing just on Ethereum, the largest blockchain ecosystem, there are many hundreds of DAOs today which collectively have more than 20 billion US Dollars locked up in their systems.[7] Collective decisions made in DAOs can include things like how to allocate shared resources, what changes to make to the project's code base, what values to set key parameters to, what business partnerships the project should enter into, and more. Early on in its life, a web3 project typically specifies which decisions are made collectively through the DAO and how. This "constitution" is written into the project's code, and collective decisions are then made via votes of the tokenholders, with the results of the votes automatically executing upon completion.

Low participation is a major challenge to DAO governance. In addition to threatening the intended egalitarian ethos of web3, low participation also poses security risks and can lead to elite capture (Qin et al. 2021). Analyzing thousands of wallets in three major decentralized finance companies, Fritsch and Wattenhofer (2022) find that less than 15% of circulating tokens participate in governance—in contrast to approximately 70% participation in traditional shareholder meetings. Messias and Loiseau (2023) find that it is unduly expensive for small token holders to vote, given high transaction fees for processing on-chain votes. Based on an analysis of token distribution and voter turnout of nine decentralized finance protocols, Barbereau et al. (2022) conclude that decentralized governance systems today resemble "timocracies" in which only property owners participate in governance.

experiments where the platforms and applications built on top of this blockchain-powered "web3" internet are trying to democratically distribute governance power among users, often via token voting. To briefly define, a blockchain is essentially a digital ledger, or record book of online transactions, that is secured by cryptography. The name stems from information entries being recorded on "blocks," and as new blocks are added to past blocks this creates a "chain" of public online records. Blockchains are decentralized because instead of one entity (e.g., AWS, Google cloud, or other corporations/ service providers) controlling this online record, many computers distributed around the globe work together to verify and store transactions on blocks.

[7] This figure is based on the DeFI Llama dashboard, accessed on Aug 22, 2023. `https://defillama.com/chain/Ethereum`.

This widespread problem of low participation in online governance has sparked a range of interesting experiments aiming to encourage more people to participate in governance in web3—in this paper, we focus on the most ambitious and largest-scale of these experiments to date.

### 2.2 Setting for our Experiment: Optimism

Optimism is an important "layer-2" blockchain built on top of the popular Ethereum blockchain that aims to make Ethereum transactions faster and more cost-effective. The Optimism Collective, as it's known, is governed by a DAO structured as a bicameral legislature that decides how their large treasury, which currently holds roughly $4 billion USD worth of tokens, is spent.[8] The Collective consists of two houses: the Token House, which aims to represent business interests through token-based voting, and the Citizens' House, which is based on online identities rather than tokens and aims to represent the broader community of Optimism stakeholders who may not be able to afford large token stakes. Anyone can purchase Optimism's OP tokens on the open market and use those OP tokens to vote in the token house. Token holders vote on topics such as protocol upgrades, treasury allocations, or code of conduct violations.[9] In an effort to publicly commit to rewarding what they consider positive contributions to the community, Optimism has earmarked 20% of its budget for public goods funding based on peer attestations, and a separate 19% for sequential airdrop rewards aimed at incentivizing prosocial behavior.

### 2.3 Details of Optimism's Airdrop 2

On February 9th, 2023, Optimism distributed 11.7 million OP tokens worth a total of approximately 28 million US Dollars to over 300,000 unique Optimism addresses via an airdrop. In the context of web3, an airdrop refers to a project directly distributing tokens to user wallets, with no action required by the user. Airdrops are typically sent to wallets of users who complete certain tasks or meet specific criteria, and might be used for marketing purposes to create interest and awareness about a new project, or to incentivize certain behaviors such as interacting with a project's tools or debugging early code. While multiple web3 platforms have experimented with airdropping

[8]This figure comes from the DeepDAO dashboard, accessed on August 23rd, 2023. See `https://deepdao.io/organizations`.

[9]See additional details in Optimism's voting manual here: `https://github.com/ethereum-optimism/OPerating-manual/blob/main/manual.md`

**Table 1 – Airdrop 2 Criteria & Allocations**

| Reward Type | Criteria | # of Qualifying Addreesses | Reward Allocation | Total OP Disbursed |
|---|---|---|---|---|
| **Governance Delegation Reward** | Has had ≥ 2,000 total Tokens x Days[(1)] | 57,204 | $^{0.42}/_{365}$ OP per Tokens x Day, max 5,000 OP per address | 6.8 million OP |
| **Gas Usage Reward** | Spent ≥ the average cost of one Ethereum transaction ($6.10) on Optimism | 280,057 | 80% of gas fees rebated in OP, up to $500 of gas fees rebated per address | 2.5 million OP |

(1) OP Delegated x Days = Cumulative Sum of OP Delegated per Day (i.e. 20 OP delegated for 100 days: 20 x 100 = 2,000 Delegated x Days).

rewards to users, there are unique design choices behind Optimism's airdrop 2 which we leverage to empirically study. Specifically, the fact that the eligibility criteria was unknown to participants ex ante, the focus on rewarding governance participation, and the large amount of individual-level, time-stamped, and publicly available data for individuals both above and below the eligibility threshold offer an unprecedented opportunity to estimate the causal effects of rewards on governance participation at scale. In our analysis, we consider governance participation to include both voting with tokens, or delegating tokens for a representative to vote on one's behalf.

The precise eligibility criteria and reward formulas for airdrop 2 are displayed in Table 1. As the table shows, the most important factor for receiving rewards is "delegation." Optimism encourages users to delegate OP tokens to fellow community members who have explicitly volunteered to play an active role in governance—similar to proxy voting in corporate governance, and akin to voting for a representative. When a user delegates their tokens to another user, the delegate gains the voting power of the delegator's tokens and can use them to vote when votes occur in the Token House. If a user chooses to vote themselves rather than select a representative to vote on their behalf, Optimism still requires the user to self-delegate their tokens. Hence, the act of delegating is a necessary component of participating in Optimism's governance. Delegation in Optimism is not costly other than the cognitive effort of deciding to delegate (and to whom), along with some standard transaction fees which we describe in greater detail in the next paragraph. Moreover, users retain 100% ownership of their delegated tokens and can use them however they want, including selling or un-delegating tokens at any point.

As the second row of Table 1 indicates, users also received reimbursements for gas fees in addition to the rewards for delegation. Gas fees are transaction costs associated with using the Optimism blockchain, for example sending or receiving tokens. These fees can vary and will increase when

many people are using the underlying blockchain network at the same time. Gas fees essentially ensure that the network runs smoothly and that user transactions get processed. Someone who actively uses the network will accrue higher gas fees over time, which is likely why Optimism included total gas usage in the reward function for Airdrop 2.

Finally, there were 4 extra attributes that allowed addresses to earn multiplier bonuses on top of the baseline criteria in Table 1. For example, if an address spent $\geq$ \$20 on gas fees, meaning these addresses did a larger sum of activity on Optimism, they received a 2x (100%) bonus to their overall airdrop reward, or if an address had $\geq$ 54,367 total 'OP Delegated x Days,' indicating a larger governance commitment, the qualifying addresses received a 1.1x (10%) bonus. These extra attributes are important because they allow us to obtain variation in reward sizes for users with similar levels of pre-treatment governance activity.

Importantly, the eligibility criteria was announced retroactively and in conjunction with the "surprise" retroactive reward delivery itself. Users knew an airdrop was coming, given that Optimism had publicly announced their decision to allocate 19% of the token supply to a series of airdrops and only 5% was spent on airdrop 1. People did not know, however, when the airdrop would occur, nor did they know how rewards would be allocated and what behaviors would be rewarded[10].

### 2.4 Theoretical Expectations

Optimism's airdrop 2 is an interesting intervention because it rewards people for participating by giving them a greater stake in the future of the community (i.e., influence on community governance policies), and because it promises more rewards in the future. Neither of these features guarantee that the intervention will increase participation, though. Critically, users can at any time sell the awarded tokens on the open market. Hence, it's plausible the intervention could have no effect or even a negative effect on participation if it leads people to exit the project.

There are two kinds of effects worth investigating with this in mind. First, if people believe OP tokens might be worth more on the open market in the future than they are today, *and* they believe Optimism's promises to offer more governance-related airdrops in the future, then the

[10]On this point, it's worth noting that Airdrop 1 rewarded completely different behaviors (none of which specifically targeted governance participation) and used a different reward function based on binary characteristics.

announcement of airdrop 2 could induce more delegation among people who haven't previously delegated. On the other hand, if Optimism's promises aren't seen as credible, or people do not believe in the future of the community or the value of the token, then the announcement will have no effect. Qualitative evidence from Discord and Twitter (both popular social media platforms for Optimism community members) confirms strong anticipation for airdrop 3, with users already commenting *"what we do for airdrop 3 lol"* on Optimism's Feb 9th tweet announcing the release of airdrop 2. Users posted many variations of this the following day on Feb 10 as well, for example *"when airdrop 3?,"* on Discord or *"waiting for next airdrop3"* on Twitter. One user commented the following on Feb 10: *"I'm willing to repeat it for as long as it takes. I support the project and thank you for round 2 airdrop. Looking forward to further development of the ecosystem and new opportunities to profit."*

Second, among people who directly received rewards in airdrop 2, this increased stake in the future of the community could lead them to want to participate more. If so, then people who receive larger rewards should, on average, participate more. On the other hand, there is nothing to stop those who receive rewards from selling them instantly.

We turn now to quantitatively evaluating these two kinds of effects.

## 3 Results: Rewards Increase Participation

As discussed above, on February 9th, 2023, Optimism revealed that addresses were receiving rewards based in part on their past rate of participation in the governance system, as measured by the number of OP tokens they delegated in the Optimism system, multiplied by the length of time the tokens remained delegated. While we first examine effects on token-day delegations, we later measure participation in terms of voting.

If people paid sufficient attention to learn about this airdrop, *and* believed that Optimism would reward governance participation in future airdrops, *and* valued the opportunity to obtain more OP tokens in the future more than the chance to sell the OP tokens immediately, then this announcement might cause people to delegate more tokens.

Figure 1 shows the jump in new delegations after the announcement of airdrop 2. The left panel plots the total amount of new OP tokens (logged) that were delegated each day. The day of the

**Figure 1 – Increase in Delegation After Announcement of Rewards Scheme on Feb 9, 2023.** Plots show daily total flows in new delegations by OP tokenholders, measured in logs.

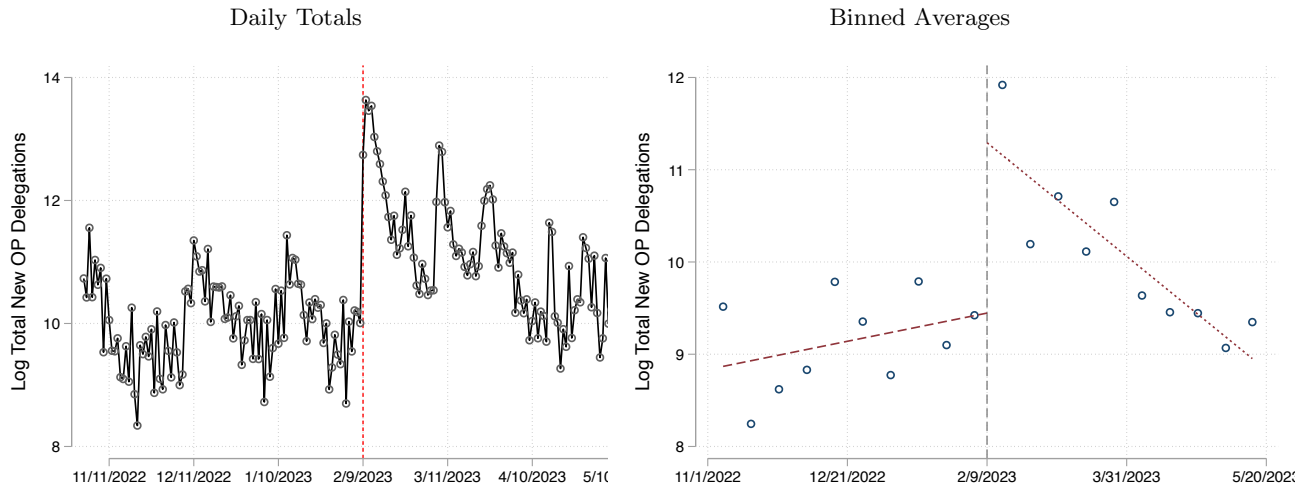


**Table 2 – Effect of Airdrop Announcement on Subsequent Delegations: Regression Discontinuity Design**

| | Total Delegated | # Wallets Delegating | Total First Time | Total Non-First |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| RD_Estimate | 289351.84 | 10338.28 | 47848.69 | 241701.13 |
| | (68192.26) | (2358.95) | (10174.45) | (59485.74) |
| Observations | 199 | 199 | 199 | 199 |

Estimates and standard errors from automated RD procedure from Calonico et al.

announcement is indicated by the red vertical line in the middle of the plot. As the plot shows, there is a noticeable spike in delegation activity in the immediate aftermath of the announcement, especially on the first two days, with new delegations remaining elevated for roughly the next week.

The daily data is inherently bouncy, so the right panel of the plot presents binned averages that pool across days to reduce some of the noise. Each point is an equal-sample-sized average of many underyling points. Lines of best fit are estimated to either side of the announcement date. As we can see, there is a clear, discontinuous jump in new delegations on the day of the announcement and following.

Exactly how large is this discontinuous jump? In column 1 of Table 2, we estimate the size of the discontinuity using a standard regression discontinuity estimator (Calonico, Cattaneo, and

Titiunik 2014). Using this approach, we estimate an increase of roughly 290,000 OP tokens on the day of the announcement. Clearly this is an underestimate of the total effect of the announcement, since much of the new delegations came in on subsequent days, but it does suggest the increase is meaningful. Prior to the announcement, the average total OP tokens newly delegated per day is a bit less than 12,000, so the 290,000 increase is sizable compared to this baseline.

In column 2, we re-estimate this effect on the number of unique wallets that make delegations. We estimate that the announcement leads to roughly 10,000 wallets making new delegations on the first day–indicating that the effect on total delegations is not entirely driven by a small number of addresses delegating a large amount of tokens.

As we discussed in section 2.3, addresses that were already delegating their OP tokens had their newly airdropped tokens automatically delegated, differing from airdrop 1 in which token rewards needed to be claimed by users. These addresses had the option to pull these delegations back and to sell the tokens at any time, but given behavioral frictions, it is not surprising to find that they mostly chose to leave the delegations in place.

Accordingly, we reestimate our effects separately for addresses that have never delegated before ("First Time" delegators) and addresses that have ("Non-First Time" delegators). As column 3 in Table 2 shows, we find a large increase of approximately 48,000 OP tokens in first-time delegations. The level of increase is larger for non-first-time delegations, as column 4 shows, but the results are roughly the same in terms of proportional increases from their previous baselines, because daily first-time delegations are smaller, on average. These results suggest that the overall effect on delegations is not only driven by user inertia for addresses whose rewards are auto-delegated.

### 3.1 Rewards to Addresses Increase Their Participation

Thus far, we have focused on the overall effect that the airdrop 2 announcement had on governance participation, because a major logic for the Optimism rewards scheme is to generate expectations of future rewards. However, this is not the only purpose—the rewards scheme also gives tokens directly to addresses in reflection of their past behavior. To what extent do these addresses commit these new tokens to governance?

Estimating the effect of receiving these rewards is complicated by a very direct problem of statistical confounding. Addresses received rewards in part as a function of their previous delegation

behavior. An address that was already participating in governance is probably likely to continue participating in governance, and is also guaranteed to receive a larger reward size. If we do not directly address the way that rewards were assigned to addresses, there is likely to be a strong correlation between governance participation and reward size that does not reflect a causal effect of receiving the reward.

Although we have no perfect way to mitigate this problem, our strategy is to estimate equations of the form

$$Delegation_{it} = \beta_0 + \beta_1 f(Log\ Rewards_{it}) + \beta_2 Delegation_{i,t-1} + X_{it} + \epsilon_{it}, \quad (1)$$

where *Delegation* measures the total delegate-days committed by address $i$ during the treatment period ($t$) or the pre-treatment period ($t-1$). By controlling for previous total delegations, we block this part of the rewards formula from the analysis. Residual variation in the size of the rewards comes from addresses that had roughly similar rates of delegation but which received larger or smaller rewards because of the amount of gas refunds and other awards bonuses they received (see Section 2.3 for the discussion of the rewards formula).[11] To further pin down the source of variation, we can also control directly for the four bonus categories, so that only variation from gas usage is used; we pursue this strategy in the main Appendix table.

Figure 2 presents our preferred version of the overall results, using the binscatter package in Stata. Each point in the plot represents an equal-sample-sized binned average of many underlying observations. The data has been residualized by pre-treatment total delegate-days, as in the lagged-DV design represented in equation 1. The plot shows a quadratic line of best fit estimated on the underlying data.

As the figure shows, there is a strong positive conditional relationship between the size of the rewards and the amount of delegation after receiving the rewards. Because the figure is residualized by pre-treatment delegation totals, the positive relationship does not simply reflect that addresses that were already delegating and hence received larger rewards automatically delegated those rewards. Indeed, another way to see this is to re-do the plot excluding all addresses that delegated

[11] The data on delegations both before and after the airdrop includes extreme outliers reflecting the delegations of very large tokenholders. These outliers cause challenges for our OLS estimates, typically leading to extreme over-estimates of the effect of rewards. To be conservative, we trim the data to exclude addresses above the 99th percentile in either pre-treatment or post-treatment delegations.

**Figure 2 – Larger rewards lead to more delegation, on average.** Binned averages of log rewards and total delegate-days; data residualized by lagged total delegate-days to approximate lagged-DV design.

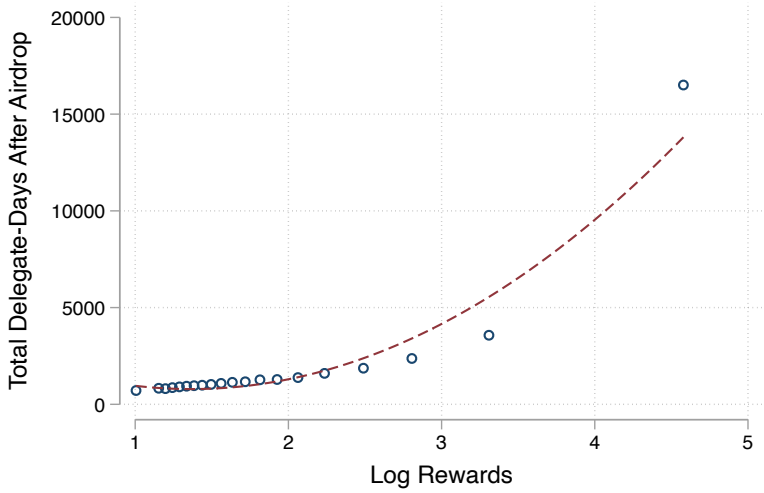


any OP tokens prior to the airdrop. The resulting plot, available in the Appendix, is extremely similar. In the Appendix, we also offer formal estimates of this relationship.

This strong positive relationship does not guarantee that the rewards directly cause the delegation behavior. Addresses that received larger rewards, conditional on pre-treatment delegation behavior, likely used Optimism for more transactions prior to the airdrop (and thus accrued more gas fees and therefore more gas refunds in the airdrop). It could be that heavier users of the Optimism system are more interested in governance and would have participated in governance even in the absence of the airdrop. However, in conjunction with the regression discontinuity results in the previous section, it seems plausible to suspect that many addresses that received rewards put them into the governance system because they learned about the promise of delegation by receiving the reward.

### 3.2 Rewards and Positive Governance Participation

As a general matter, a major downside to paying for participation in governance is that there is no guarantee the participation will be thoughtful or beneficial. In web3 this could be an especially big challenge, because adversarial behavior is common and people can code bots to try to harvest

**Figure 3 – Larger rewards lead to more delegation, and especially to more delegation to other addresses.** Binned averages of log rewards and total delegate-days; data residualized by lagged total delegate-days to approximate lagged-DV design.

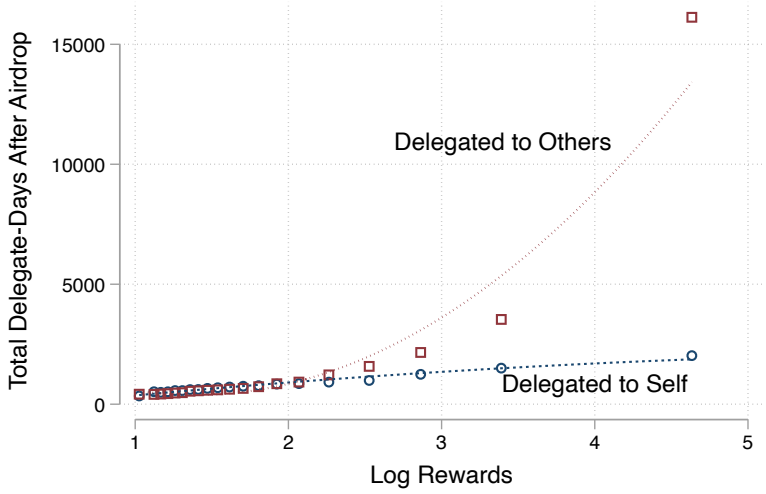


rewards. It is not really possible to define or measure "good" participation. However, we can look at several intermediate outcomes.

The first thing we examine is to whom addresses delegate their tokens. When users delegate their tokens in order to participate in Optimism's governance system, they have the choice of delegating their tokens to themselves, in which case they can directly vote on proposals, or delegating their tokens to someone else who they prefer to vote in their stead. Optimism encourages people to delegate their tokens to experts who have the time and skills necessary to study upcoming proposals before deciding how to vote. However, airdrop 2's reward system did not directly incentivize delegating to others instead of to selves, so we can examine whether users voluntarily chose to delegate to others more often.

Figure 3 again shows the conditional correlation between reward size and delegation behavior after residualizing on pre-treatment delegation behavior. This time, though, we fit separate quadratic curves for addresses that delegated their tokens to others vs. addresses that delegated to themselves. As the plot shows, where the reward size grows large, addresses delegate to others much more often than to themselves.

In the Appendix, we present formal estimates of this relationship, confirming that the relationship between reward size and delegation size is much stronger for delegations to others than for

delegations to selves. Moreover, we also show that the effect of rewards on delegation is larger for addresses with smaller OP token balances prior to the airdrop—indicating that the rewards scheme is bringing smaller tokenholders into the governance system.

### 3.3 New Delegations Are Dispersed Across More Delegates

Clearly, airdrop 2 led people to delegate more tokens to other delegates. But whom did they delegate to? If the decision to delegate was largely mindless and purely focused on harvesting future rewards, we might expect people to delegate to the most popular current delegates as the path of least cognitive resistance. Accordingly, we next examine the distribution of delegations before and after the announcement of the airdrop. Focusing on the aggregate distribution of delegations made one month before and one month after the announcement, we find that there is much more dispersion after the airdrop than before. For example, one way to measure concentration is HHI, a popular measure of market concentration used in economics.[12] By this measure, we find that concentration is more than twice as large prior to the announcement than after (2,277 prior to treatment; 710 afterwards). We can also see this increased dispersion in the number of delegations, not just in the total number of tokens delegated; the average number of delegations a delegate received before treatment was 46, while after it was 38. In sum, the airdrop increased delegations to others, and did so without encouraging clumping on the most popular delegates—instead, it encouraged a broader range of delegation to more different delegates.

### 3.4 Rewards Increase Voting Rate

Our final analysis looks directly at the act of voting. Airdrop 2 only explicitly rewarded delegation, not voting. Delegation is a necessary step to voting, but delegated tokens can still choose not to vote. In Figure 4, we show the conditional relationship between reward size and a count of the number of votes cast by an address during the post-treatment period—giving addresses credit for voting if they either voted directly on a proposal, or if their delegate voted for them. Unfortunately, we are only able to access voting data during the post-treatment period, so we cannot include lagged voting as a control. Instead, following previous analyses, the figure controls for lagged delegation behavior as a proxy.

[12] See `https://www.justice.gov/atr/herfindahl-hirschman-index`.

**Figure 4 – Larger rewards lead to more voting.** Binned averages of log rewards and number of votes; data residualized by lagged total delegate-days to approximate lagged-DV design.

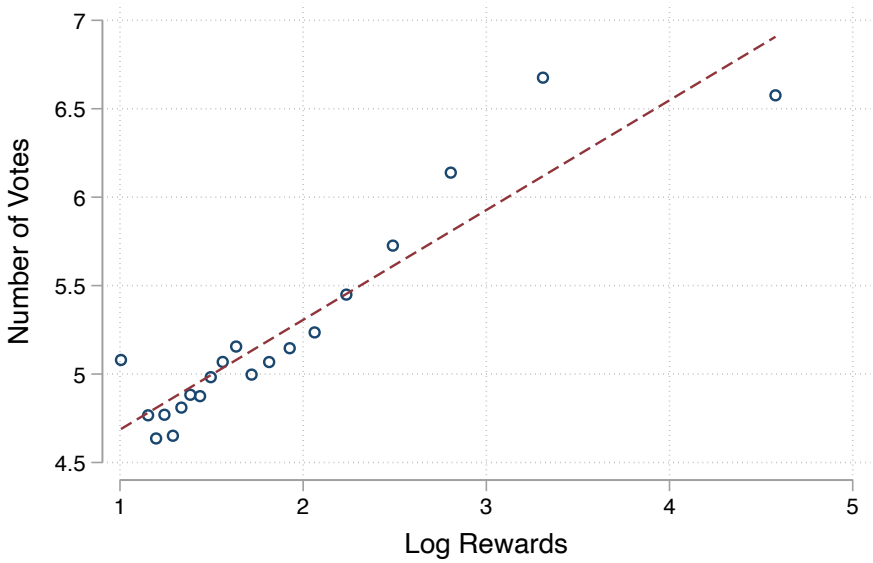


As the figure shows, we a find strong relationship between reward size and voting. The Appendix provides formal estimates to complement this figure.

## 4 Conclusion

Increasing political participation is an important goal of democratic systems in the physical world and, increasingly, in the online world as well. Why and when people decide to participate in democracy is difficult to understand—instrumental incentives to vote are often small, because individuals have little chance of affecting outcomes in most elections of any reasonable size. As a result, there has been a lot of interest in figuring out what kinds of incentives *do* lead people to vote, beyond purely instrumental incentives regarding the outcome of the vote itself. Experiments in the physical world suggest that directly compensating people to vote have typically modest effects, but the interventions that have been empirically testable to date have all been one-off payments. The intervention we study in this paper aims to induce deeper, longer-run engagement in governance by providing rewards in the forms of tokens that provide an ongoing stake in the community, and offering them in a sequence over time. As we have shown, the announcement of rewards for participation increased participation by people hoping to earn future rewards, and people who

gained a larger stake in the project by receiving rewards directly responded by participating more. The reward scheme had larger effects for smaller tokenholders and encouraged a broader dispersion of delegation, suggesting a broadening of democracy.

There are important limitations to our analyses. First, although we are able to observe roughly five months of data subsequent to the intervention, the effects we document could be temporary and might fade out in the longer run. This is standard in experimental work on political behavior (for example see discussion in Broockman and Green 2014), though, and our ability to measure behavior five months after treatment is reassuring in this regard. Second, there is no direct way to measure the quality of participation, so we cannot completely rule out that some or all of the delegation behavior we observe is driven purely by rewards harvesting. This is a general challenge in the study of compensating people for political participation. We actually have the advantage of being able to measure some partial indicators of how meaningful the participation induced by the intervention is, because we can see which delegates people delegate to and whether they vote. Nonetheless, there is no way to rule out that the delegation is uninformed or that the delegates themselves make bad decisions when voting. This is a fundamental obstacle in the study of voting. Finally, we have no way to establish the unique identities of the wallets we study. It is possible that one person could possess many wallets, and that we could as a result overestimate the number of unique people participating in Optimism's governance. Because Optimism's reward structure was specifically designed to eliminate any incentive to split tokens across wallets—rewards were based purely on how many tokens the wallet had previously delegated, so one wallet delegating 10 tokens would receive the exact same total reward as two wallets each delegating 5 tokens, for example—this is unlikely to be a major problem. It is also unlikely that the reward limit of 5,000 token-days would be gamed by splitting for example 10,000 token-days into two wallets of 5,000 token-days, given that the reward limit was unknown prior to the airdrop implementation. However, this could be a challenge to future research in other settings, in which case deeper efforts to estimate the probability that multiple addresses map to the same user may be required.

Our results raise a number of open questions for governance in both the physical world and the digital world. The impact of Optimism's reward scheme on participation seems to have been driven both by creating a credible promise to reward participation in the long run, not just the short run, and by denominating rewards in terms of a token that gives people a stake in the future

of the project. This is consistent with evidence that other forms of long-term economic stakes like homeownership (Hall and Yoder 2022) or universal basic income payment programs (Loeffler 2023) increase turnout in the physical world. In a world where voters are often present-biased (Hill 2020), the combination of direct immediate rewards plus the promise of future rewards could help to increase participation, especially in local elections where turnout is much lower in the US. It seems valuable for political scientists to study how giving people a durable stake in their local community might increase participation—including ideas like encouraging participation in financial markets, as in Jha and Shayo (2019), encouraging homeownership, or other similar interventions.

In the specific case of Optimism, a key question is how long the sequence of rewards can go on and how many different behaviors can be rewarded productively over time. If delegation continues to be a focus of future rewards, adversarial rewards harvesters may begin to crowd out pro-social behavior, delegating mindlessly and eroding the governance system. To avoid this, the project can seek to stay one step ahead, changing what kinds of pro-social behavior are rewarded, but there are only so many observable components to governance participation. A reputation system based on early airdrop recipients, who offer a kind of ground truth for pro-social behavior, may be another way to define future rewards.

For other online communities, there are a number of unique aspects to Optimism's reward scheme that deserve study. Clearly, the announcement of airdrop 2 was able to move people's beliefs—implying both that people (a) paid attention to the announcement, and (b) found Optimism's commitment to future airdrops credible, so that taking actions today could pay off in the future. If people do not pay attention to the announcements or do not believe in the project's promise to deliver future rewards, announcements would presumably have no effect. Finally, at the time of the airdrop, the market value of an OP token was relatively low. This may have been important in preventing people from immediately cashing out the rewards and instead remaining committed to the project.

At a more general level we hope that our paper showcases the potential for studying deep governance questions empirically using web3. This approach has three distinct advantages. First, web3 projects are experimenting with a range of different ways to design governance systems and incentives that we haven't observed in the physical world. Second, these experiments are going on with many thousands of users across the world, providing remarkable scale at which to analyze

impacts. And third, the data necessary to analyze these experiments is both highly granular and completely public, allowing us to observe the full range of outcomes necessary to study the effects of governance interventions. Using this laboratory, we have studied a highly unique policy that stimulated broad participation using token-based incentives. In the future, there are many other interventions that can be studied in a similar manner.

## References


Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020. "The Welfare Effects of Social Media." *American Economic Review* 110(3): 629–676.

Barbereau, Tom, Reilly Smethurst, Orestis Papageorgiou, Johannes Sedlmeir, and Gilbert Fridgen. 2022. "Decentralised Finance's Timocratic Governance: The Distribution and Exercise of Tokenised Voting Rights." *Technology in Society* 73.

Blais, André, and Agniezka Dobrzynska. 1998. "Turnout in Electoral Democracies." *European Journal of Political Research* 33(2): 239–261.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. "A 61-Million-Person Experiment in Social Influence and Political Mobilization." *Nature* 489(7415): 295–298.

Broockman, David E., and Donald P. Green. 2014. "Do Online Advertisements Increase Political Candidates' Name Recognition or Favorability? Evidence from Randomized Field Experiments." *Political Behavior* 36: 263–289.

Bryan, Christopher J., Gregory M. Walton, Todd Rogers, and Carol S. Dweck. 2011. "Motivating Voter Turnout by Invoking the Self." *Proceedings of the National Academy of Sciences* 108(31): 12653–12656.

Calonico, Sebastian, Mattias D. Cattaneo, and Rocío Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82(6): 2295–2326.

Fowler, Anthony. 2013. "Electoral and Policy Consequences of Voter Turnout: Evidence from Compulsory Voting in Australia." *Quarterly Journal of Political Science* 8(2): 159–182.

Fritsch, Robin, Müller Marino, and Roger Wattenhofer. 2022. "Analyzing Voting Power in Decentralized Governance: Who controls DAOs?" *Working paper, preprint available at* `https://arxiv.org/pdf/2204.01176.pdf` .

Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment." *American Political Science Review* 102(1): 33–48.

Gonzales, Mariella, Gianmarco León-Ciliotta, and Luis R. Martínez. 2022. "How Effective Are Monetary Incentives to Vote? Evidence from a Nationwide Policy." *American Economic Journal: Applied Economics* 14(January): 293–326.

Green, Donald P., and Alan S. Gerber. 2015. *Get Out the Vote: How to Increase Voter Turnout.* Brookings Institution Press.

Guess, Andrew M., Neil Malhotra, Jennifer Pan, Pablo Barberá, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow et al. 2023. "How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election Campaign?" *Science* 381(6656): 398–404.

Hall, Andrew B., and Jesse Yoder. 2022. "Does Homeownership Influence Political Behavior? Evidence from Administrative Data." *The Journal of Politics* 84(1): 351–366.

Highton, Benjamin, and Raymond E. Wolfinger. 2001. "The Political Implications of Higher Turnout." *British Journal of Political Science* 31(1): 179–223.

Hill, Seth J. 2020. "Following Through on an Intention to Vote: Present Bias and Turnout." *Political Science Research and Methods* 8(4): 803–810.

Jackman, Robert W. 1987. "Political Institutions and Voter Turnout in the Industrial Democracies." *The American Political Science Review* 81(2): 405–423.

Jha, Saumitra, and Moses Shayo. 2019. "Valuing Peace: The Effects of Financial Market Exposure on Votes and Political Attitudes." *Econometrica* 87(5): 1561–1588.

La Raja, Raymond J., and Brian F. Schaffner. 2022. "A Cash Lottery Increases Voter Turnout." *Plos one* 17(6): e0268640.

Lijphart, Arend. 1997. "Unequal Participation: Democracy's Unresolved Dilemma." *The American Political Science Review* 91(1): 1–14.

Loeffler, Hannah. 2023. "Does a Universal Basic Income Affect Voter Turnout? Evidence from Alaska." *Political Science Research and Methods* 11(3): 521–536.

Marble, William, and Clayton Nall. 2021. "Where Self-interest Trumps Ideology: Liberal Homeowners and Local Opposition to Housing Development." *The Journal of Politics* 83(4): 1747–1763.

Messias, Johnnatan, Pahari Vabuk Chandrasekaran Balakrishnan Gummadi Krishna P., and Patrick Loiseau. 2023. "Understanding Blockchain Governance: Analyzing Decentralized Voting to Amend DeFi Smart Contracts." *Under review, preprint available at* `https://arxiv.org/pdf/2305.17655.pdf` .

Panagopoulos, Costas. 2013. "Extrinsic Rewards, Intrinsic Motivation and Voting." *The Journal of Politics* 75(1): 266–280.

Qin, Kaihua, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. 2021. "Attacking the DeFi Ecosystem with Flash Loans for Fun and Profit." *Working paper available at* `https://www.doc.ic.ac.uk/~livshits/papers/pdf/fc21b.pdf` .

Shineman, Victoria Anne. 2018. "If You Mobilize Them, They Will Become Informed: Experimental Evidence that Information Acquisition is Endogenous to Costs and Incentives to Participate." *British Journal of Political Science* 48(1): 189–211.

Verba, Sidney, and Norman H. Nie. 1972. *Participation in America: Political Democracy and Social Equality*. University of Chicago Press.

Yoder, Jesse. 2020. "Does Property Wwnership Lead to Participation in Local Politics? Evidence from Property Records and Meeting Minutes." *American Political Science Review* 114(4): 1213–1229.

# Appendix

*Intended for online publication only.*

## Formal Estimates of Direct Rewards Effect

In this section, we provide a table of formal estimates related to the analysis presented in Figure 2 and Figure 3.

**Table A.1 – Effects of Rewards on Delegation.**

| | Total Delegate-Days | | Delegated (Binary) | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Log Total Rewards | 1308.782 | 1549.897 | 0.022 | 0.041 |
| | (34.260) | (38.590) | (0.001) | (0.001) |
| Prior Delegate-Days | -0.005 | -0.050 | | |
| | (0.019) | (0.019) | | |
| Bonus Category 1 | 11015.390 | 10880.320 | 0.106 | 0.080 |
| | (153.003) | (151.926) | (0.002) | (0.002) |
| Bonus Category 2 | 17686.055 | 16909.248 | 0.281 | 0.187 |
| | (305.016) | (311.506) | (0.006) | (0.005) |
| Bonus Category 3 | 242.963 | 67.954 | 0.009 | -0.015 |
| | (30.787) | (33.564) | (0.002) | (0.002) |
| Bonus Category 4 | -1042.466 | -1092.449 | -0.110 | -0.105 |
| | (53.636) | (58.019) | (0.003) | (0.003) |
| Prior Token Balance | | -0.000 | | -0.000 |
| | | (0.000) | | (0.000) |
| Prior Delegated (Binary) | | | 0.559 | 0.497 |
| | | | (0.001) | (0.001) |
| Observations | 284332 | 254997 | 284332 | 254997 |

Robust standard errors in parentheses.

## Do Users Cash Out Their Rewards?

A main reason to think that airdrops like Optimism's might not lead to more governance participation is that users can sell the token rewards they receive on the open market. As we discussed above, the price of an OP token was roughly $2.40 on the day of the airdrop 2 announcement.

Figure A.1 investigates whether we see evidence of people cashing out their rewards. Each point in the plot indicates the number of wallets with a newly zeroed-out balance on each day. As

the plot shows, during Optimism's airdrop 1, significant rates of cashing out are observed. During airdrop 2 there is again a spike in cash-outs, but the level is much lower than during airdrop 1.

**Figure A.1 – Investigating "Dumping" and Exit: Notably Fewer Wallets with 0 OP Balance at Airdrop 2 versus Airdrop 1.**

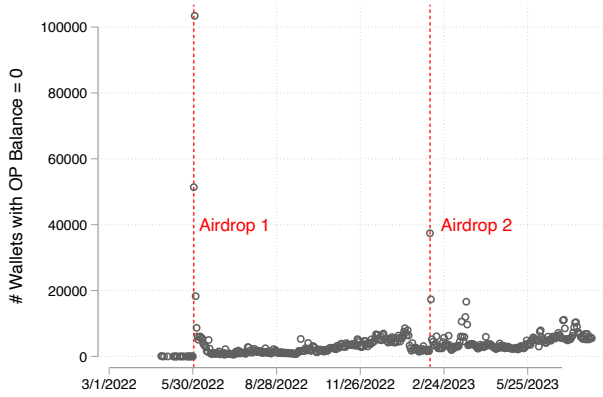


## Heterogeneity in the Effects of Rewards

In the table below, we present formal results on the variation in the direct effect of rewards on participation. In columns 1 and 2 we compare the effect of rewards on self delegation vs. delegation to others, finding that the result for delegation to others is much larger than the effect on delegation to self. In column 3 we show this another way, finding that as the size of the rewards increases, the percentage of tokens delegated to others increases noticeably. Finally, in column 4 we show that the effect of rewards on delegation is larger for smaller token holders.

**Table A.2 – Variation in Effects of Rewards on Delegation.**

| Reward Type | Criteria | # of Qualifying Addreesses | Reward Allocation | Total OP Disbursed |
|---|---|---|---|---|
| **Governance Delegation Reward** | Has had $\geq$ 2,000 total Tokens x Days[(1)] | 57,204 | $^{0.42}/_{365}$ OP per Tokens x Day, max 5,000 OP per address | 6.8 million OP |
| **Gas Usage Reward** | Spent $\geq$ the average cost of one Ethereum transaction ($6.10) on Optimism | 280,057 | 80% of gas fees rebated in OP, up to $500 of gas fees rebated per address | 2.5 million OP |

Robust standard errors in parentheses. Small holders are those whose maximum OP wallet balance prior to Feb 9, 2023 was below 1,000. Large holders are those whose balance was above 1,000.

## Effects on Voting

In this section, we present formal estimates of the direct effect of rewards on voting.

**Table A.3 – Effects of Rewards on Voting.**

| | Number of Votes | | Voted (Binary) | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Log Total Rewards | 0.621 | 0.793 | 0.024 | 0.036 |
| | (0.014) | (0.015) | (0.001) | (0.001) |
| Prior Delegate-Days | 0.000 | 0.000 | 0.000 | 0.000 |
| | (0.000) | (0.000) | (0.000) | (0.000) |
| Prior Token Balance | | -0.000 | | -0.000 |
| | | (0.000) | | (0.000) |
| Observations | 284332 | 254997 | 284332 | 254997 |

Robust standard errors in parentheses.

## Effect of Rewards Excluding Wallets That Have Previously Delegated

In the figure below, we regenerate Figure 2 excluding all wallets that had any delegations prior to the airdrop. The included wallets therefore only received rewards based on their gas fees, and not based on any participation in governance. As the figure shows, we continue to find a strong, positive relationship between rewards and subsequent participation in governance as measured through delegations.

**Figure A.2 – Larger rewards lead to more delegation, on average.** Binned averages of log rewards and total delegate-days; data residualized by lagged total delegate-days to approximate lagged-DV design. Figure excludes wallets with delegations prior to airdrop 2.