No. 84
DISCUSSION PAPER SERIES

Response to "A Comment on Vulnerability and Clientelism (2022)"

Gustavo J. Bobonis

Paul Gertler

Marco Gonzalez-Navarro

Simeon Nichter

This paper responds to:
Ma, Hai, Sébastien Montpetit, and Ardyn Nordstrom. 2023. A Comment on "Vulnerability and Clientelism" (2022). I4R Discussion Paper Series No. 83. Institute for Replication.

November 2023

INSTITUTE for REPLICATION

I4R Discussion Paper Series

I4R DP No. 84

Response to “A Comment on Vulnerability and Clientelism (2022)”

Gustavo J. Bobonis[1], Paul Gertler[2], Marco Gonzalez-Navarro[2], Simeon Nichter[3]

[1]University of Toronto/Canada
[2]University of California, Berkeley/USA
[3]University of California, San Diego/USA

November 2023

ISSN: 2752-1931

Response to "A Comment on Vulnerability and Clientelism (2022)"

Gustavo J. Bobonis, Paul Gertler,
Marco Gonzalez-Navarro, and Simeon Nichter

October 2023

Bobonis, Gertler, Gonzalez-Navarro, and Nichter (2022) conducts a randomized control trial in rural Northeast Brazil designed to reduce the vulnerability of sampled households. In this development intervention, we constructed residential water cisterns across 425 neighborhood clusters in 40 municipalities, and examine effects using a longitudinal panel survey and electoral data at the precinct level. Ma, Monpetit, and Nordstrom's (2023) comment confirms the reproducibility of our results. Moreover, their comment does not challenge any of our article's primary findings: the cisterns treatment significantly reduced citizens' vulnerability (Table 2), it decreased citizens' requests for private goods from politicians (Table 3), and it significantly decreased votes for incumbent mayors (Table 4).

The comment by Ma, Monpetit, and Nordstrom (2023) discusses three aspects of robustness: (1) the matching of individuals in the panel over time, (2) how clientelist relationships are defined, and (3) the choice of historical rainfall period. With regards to the first aspect, the comment reports some age inconsistencies across waves for a relatively small subsample, even though it states that results remain "stable in terms of both magnitude and statistical significance" when excluding these observations. As discussed below, our longitudinal rostering procedure accurately identifies individuals across survey waves, though some minor measurement error exists in reported ages.

With regards to the second aspect, the comment challenges Section VI of our article, which presents additional heterogeneity analyses in Table 5 to explore the role of clientelism in our primary results. More specifically, the comment argues that those results are not robust to a more restrictive coding of the binary clientelism marker employed to test heterogeneity. Contrary to their critique, we show that analyses in Table 5 of our article are indeed robust to a more restrictive coding. With regards to the third aspect, their comment indicates that halving the window of historical data used to normalize rainfall affects only a single, ancillary result: the cistern treatment's impact on one of three well-being measures we examine (Column 3 in Table 2). Since Ma, Monpetit, and Nordstrom (2023) indicate that "the overall message remains the same" — and it is not obvious that their approach is preferable — we do not discuss the third aspect below.

1 Matching of individuals over time

Ma, Monpetit and Nordstrom (2023) report that some age inconsistencies exist across waves, though they indicate that results remain "stable in terms of both magnitude and statistical significance" when excluding these observations. More specifically, the comment reports that: "Given the timing between wave 1 and wave 2, we would expect individuals to be between 0 and 2 years older at wave 2 than they were at wave 1. We identified that 9.2% of the sample is either younger at wave 2, or more than 2 years older." We thoroughly investigated the purported inconsistencies using the non-anonymized dataset, and find that our longitudinal rostering procedure correctly matches individuals across waves. We emphasize that the discussion below focuses on a small subset of our study's sample, which includes 2,680 individuals in the 2012 wave and 1,944 individuals in the 2013 wave.

To investigate the purported inconsistencies, we conducted a name match to determine whether the identical individuals answered each wave of our survey. To reduce false negative matches, we stripped accent marks and capitalization from all waves. We then employed Stata's matchit algorithm. Using this procedure, 97% of 248 individuals that Ma, Monpetit and Nordstrom (2023) purport to having age discrepancies actually have *exact* name matches across waves (i.e., with a 100% matching threshold). For the remaining 10 observations, we conducted a manual inspection, which suggested that seven of these non-matches were due to mispelling, omission of a suffix (Junior), or a different household member taking the survey.

As such, the non-anonymized dataset — which Ma, Monpetit and Nordstrom (2023) could not analyze due to IRB requirements — suggests that our longitudinal rostering procedure correctly matches individuals across waves. We acknowledge that as might be expected with any survey data, some minor measurement error exists in reported ages: 2.8% of panelists have age differences across waves that are over three years greater in absolute value than expected. As noted by Ma, Monpetit and Nordstrom (2023), results are robust when excluding observations with *any* measurement error in age.

2 Alternative definitions of the clientelism marker

The primary results in Bobonis, Gertler, Gonzalez-Navarro and Nichter (2022) focus on the effects of two exogenous shocks to vulnerability (Tables 3 and 4). Next, Section VI of the article conducts heterogeneity analyses to test whether effects are stronger for citizens who are likely to be in clientelist relationships. Given that the extant literature lacks a well-established marker for whether citizens are likely to be involved in clientelist relationships, we employed a binary proxy: whether the respondent conversed at least monthly with a local politician before the 2012 electoral campaign began. The second component of Ma, Monpetit and Nordstrom's (2023) comment examines whether heterogeneity analyses of the citizen requests outcome is robust when coding this clientelism marker more restrictively. Contrary to their critique, we show that analyses in Table 5 of our article are indeed robust when the marker is defined more restrictively.

Ma, Monpetit and Nordstrom (2023) argue that while our heterogeneous results in Section VI hold when defining the clientelism marker as citizens with at least monthly interactions with politicians, they do not hold when defining the marker as citizens with either at least weekly or at least daily interactions with politicians. Their sensitivity analysis estimates regression models analogous to those in equation (5) of Bobonis, Gertler, Gonzalez-Navarro and Nichter (2022), in which the comment sequentially includes interactions of our exogenous vulnerability shocks with their two more restrictive definitions of the clientelism marker.

First, we emphasize that Ma, Monpetit and Nordstrom's (2023) approach has less power to detect differential effects — especially for the daily interactions category — due to the substantially smaller sample of individuals in the interacted group when employing their alternative definitions of the marker. As shown in Table 1 below, whereas 490 respondents report having at least monthly interactions with politicians (i.e., the clientelism marker in our original article), only 98 report at least daily interactions, and 304 report at least weekly interactions. And these figures are roughly halved for citizens in cisterns treatment group: whereas 238 respondents report having at least monthly interactions with politicians, only 40 report at least daily interactions, and 147 report at least weekly interactions. These numbers are so small for the daily interactions category that it is no surprise that coefficients with that marker fail to reject the null hypothesis. Correspondingly, we show below that heterogeneous results in Table 5 of our article are robust when using a weekly interactions marker, and also show that the sign and magnitude of coefficients are consistent when using a daily interactions marker.

Second, we underscore that the approach undertaken by Ma, Monpetit and Nordstrom (2023) introduces bias; once their suggested alternative clientelism marker is appropriately coded, results in Table 5 of our article are indeed robust to more restrictive definitions. Their sensitivity analysis includes a substantial subset of individuals — those with monthly interactions with politicians — in the reference category, which reduces differential effects between those with more interactions (i.e., daily and weekly) and those with few interactions (i.e., monthly or fewer). Their approach biases the estimate on the reference category upwards, and biases downward the estimate of the interaction to zero.

Before addressing this bias, we first replicate and expand on the analysis in Ma, Monpetit and Nordstrom (2023), in order to show that even with the bias that their approach introduces, we observe that point estimates are stable, albeit marginally smaller and less precisely estimated. Building on their comment's Table 6, the first column of Table 2 below shows that the cistern's treatment effect on requests is substantially greater among individuals with at least weekly interactions with politicians. The reduction in requests among citizens with this alternative clientelism marker is 7.5 percentage points (significant at the 10 percent level; p-value = 0.065).

The effects for the reference category, which now suboptimally combines individuals with monthly interactions as well as those with fewer or no interactions with politicians, implies a reduction in requests of 2.4 percentage points (significant at the 10 percent level; p-value = 0 .060). Column 2 examines heterogeneous effects of the negative rainfall shock.

A one standard deviation decrease in rainfall increases requests in the reference group — again, suboptimally including individuals with monthly interactions — by 2.2 percentage points, and there is no differential effect among the now smaller category of individuals with at least weekly interactions with politicians. The point estimate of the overall effect among citizens with this marker is stable, but given the smaller sample and the more muted comparison between groups, we should be less able to detect differential effects. Similar patterns hold in other specifications shown in the article (see Table 2, columns 3-5).

For completeness, we also replicate and expand another component of the heterogeneity analysis in Table 6 of Ma, Monpetit and Nordstrom (2023), which employs a highly restrictive marker of clientelism: daily interactions with politicians. As noted above, this definition is roughly five times more restrictive than our preferred monthly measure: only 98 respondents in our overall sample — and 40 respondents in the treatment group — report at least daily interactions (vs 490 and 238 with monthly interactions, respectively). We show in Table 3 that the ITT effect among individuals in the reference category – which suboptimally includes those with weekly, monthly, or irregular interactions — indicates a 2.7 percentage point reduction in requests (significant at the 1 percent level). The point estimates of the differential and overall effects among individuals with daily interactions are still substantial (6.6 and 9.2 percentage points), but imprecisely estimated given the drastically smaller sample. Similar patterns hold across other specifications and for rainfall using this highly restrictive proxy measure.

Thus far, we have shown — even with the bias that the comment's approach introduces — the stability of point estimates when using more restrictive definitions of the clientelism marker. Although it is the case that the differential effects estimates are marginally smaller and less precisely estimated, we observe stability in the point estimates of the overall effects among those with at least weekly or at least daily interactions. Moroever, these are significant at the 10 percent level when the more restrictive clientelism marker is sufficiently large (i.e., at least weekly interactions).

Next, we address bias in this approach by Ma, Monpetit and Nordstrom (2023), and show that upon appropriately coding the suggested alternative clientelism marker, results in Table 5 of our article are indeed robust to more restrictive definitions. Again, their sensitivity analysis includes a substantial subset of individuals — those with monthly interactions with politicians — in the reference category, which reduces differential effects between those with more interactions (i.e., daily and weekly) and those with few interactions (i.e., monthly or fewer). To eliminate the bias introduced by their approach, we run a sensitivity analysis that estimates the differential effect among respondents with highly frequent interactions — without lumping together in the reference category respondents with frequent interactions and those with no interactions at all.

To this end, Table 4 estimates models in which we allow for heterogeneous effects both for individuals who report having (i) monthly interactions with politicians, and those with (ii) at least weekly interactions. This approach addresses the bias discussed above by ensuring that the reference category remains the same as in our article — citizens who either never interact with politicians or who do so less than once a month. As shown in Column 1, we observe strong evidence of treatment effects on requests among individuals with monthly interactions with politicians (16.3 percentage points reduction; p-value = 0.001), as well as among individuals with at least weekly interactions (7.5 percentage points; p-value = 0.064). Moreover, we cannot reject the hypothesis that the effects of the cistern treatment for both groups of frequent interactors are equivalent (e.g., p-value = 0.145 in column 1). In addition, observe in the first row of Column 1 that by avoiding the inappropriate grouping of citizens with monthly interactions in the reference category, the coefficient on the cisterns treatment effect for the reference category is again small and no longer significantly different from zero (as expected).

Similar patterns hold across other specifications and for rainfall using this more appropriate approach of evaluating more restrictive clientelism markers. In sum, contrary to the critique in Ma, Monpetit and Nordstrom (2023), analyses in Table 5 of our article are indeed robust when the clientelism marker is defined more restrictively.

3 Conclusion

We appreciate the investigation of Ma, Monpetit, and Nordstrom (2023), which confirms the reproducibility of our results and does not challenge any of our paper's primary findings. In particular, their comment considers three aspects of robustness: (1) the matching of individuals in the panel over time, (2) how clientelist relationships are defined, and (3) the choice of historical rainfall period. As shown above, our study employs an accurate longitudinal rostering procedure, our heterogeneity analyses are robust to a more restrictive coding of the clientelism marker, and our findings are robust to the choice of historical rainfall period.

Table 1: Count of Individuals with Marker of Clientelist Relationship

	Clientelist Relationship	Clientelist Relationship $\times$ Treatment
At least Monthly Interactions w/ Politician	490	238
At least Weekly Interactions w/ Politician	304	147
At least Daily Interactions w/ Politician	98	40

Table 2: Citizen Requests and Heterogeneity by Clientelist Relationship (Weekly)

	Request Any Private Good			Request Any Private Good Excluding Water	Request and Receive Any Private Good
	(1)	(2)	(3)	(4)	(5)
β_1: Cisterns Treatment	-0.024*		-0.024*	-0.025**	-0.001
	(0.013)		(0.013)	(0.012)	(0.010)
β_2: Cisterns Treatment $\times$ Clientelist Relationship (W)	-0.051		-0.050	-0.014	-0.054*
	(0.041)		(0.042)	(0.038)	(0.032)
β_3: Rainfall Shock		-0.022**	-0.023**	-0.015	-0.010
		(0.010)	(0.010)	(0.009)	(0.009)
β_4: Rainfall Shock $\times$ Clientelist Relationship (W)		-0.006	-0.003	0.001	-0.012
		(0.018)	(0.018)	(0.017)	(0.014)
β_5: Clientelist Relationship (W)	0.075**	0.050**	0.075**	0.039	0.062**
	(0.031)	(0.021)	(0.032)	(0.029)	(0.025)
Effect of Cisterns Treatment for Individuals in Clientelist Relationship (Weekly):					
$\beta_1 + \beta_2$	-0.075*		-0.074*	-0.039	-0.055*
	(0.041)		(0.041)	(0.038)	(0.032)
Effect of Positive 1 SD Rainfall Shock for Individuals in Clientelist Relationship (Weekly):					
$\beta_3 + \beta_4$		-0.028	-0.026	-0.014	-0.022
		(0.018)	(0.018)	(0.016)	(0.014)
Municipality Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	4,288	4,288	4,288	4,288	4,284
Mean of Y: Treatment Group	0.149	0.149	0.149	0.120	0.089
Mean of Y: Control Group	0.177	0.177	0.177	0.146	0.094
Mean of Y: Clientelist Relationship (W) in Control Group	0.257	0.257	0.257	0.196	0.163

Notes: Outcome variable is coded 1 if respondent reported: requesting a private good (columns 1- 3)/requesting a private good excluding water (column 4)/requesting and receiving a private good (column 5) from a local politician in 2012 or 2013; 0 otherwise. Specifications employ pooled data to examine requests in either year. Cisterns treatment is coded 1 if respondent's household is in a neighborhood cluster selected for treatment; 0 otherwise. Rainfall shock is measured as the difference between rainfall in January-September of the relevant year and its historical municipal mean during identical months in 1986- 2011, divided by the municipality's historical monthly standard deviation of rainfall. Standard errors are clustered at the neighborhood level and reported in parentheses.

Table 3: Citizen Requests and Heterogeneity by Clientelist Relationship (Daily)

	Request Any Private Good			Request Any Private Good Excluding Water	Request and Receive Any Private Good
	(1)	(2)	(3)	(4)	(5)
β_1: Cisterns Treatment	-0.027**		-0.027**	-0.025**	-0.005
	(0.013)		(0.013)	(0.012)	(0.010)
β_2: Cisterns Treatment $\times$ Clientelist Relationship (D)	-0.066		-0.069	-0.028	-0.056
	(0.074)		(0.075)	(0.066)	(0.056)
β_3: Rainfall Shock		-0.025**	-0.026**	-0.015*	-0.013
		(0.010)	(0.010)	(0.009)	(0.009)
β_4: Rainfall Shock $\times$ Clientelist Relationship (D)		0.033	0.035	0.005	0.020
		(0.037)	(0.036)	(0.034)	(0.030)
β_5: Clientelist Relationship (D)	0.078	0.051	0.078	0.049	0.053
	(0.059)	(0.040)	(0.059)	(0.050)	(0.042)
Effect of Cisterns Treatment for Individuals in Clientelist Relationship (Daily):					
$\beta_1 + \beta_2$	-0.092		-0.096	-0.054	-0.061
	(0.072)		(0.073)	(0.064)	(0.055)
Effect of Positive 1 SD Rainfall Shock for Individuals in Clientelist Relationship (Daily):					
$\beta_3 + \beta_4$		0.008	0.009	-0.010	0.007
		(0.036)	(0.035)	(0.033)	(0.029)
Municipality Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	4,288	4,288	4,288	4,288	4,284
Mean of Y: Treatment Group	0.149	0.149	0.149	0.120	0.089
Mean of Y: Control Group	0.177	0.177	0.177	0.146	0.094
Mean of Y: Clientelist Relationship (D) in Control Group	0.264	0.264	0.264	0.209	0.154

Notes: Outcome variable is coded 1 if respondent reported: requesting a private good (columns 1- 3)/requesting a private good excluding water (column 4)/requesting and receiving a private good (column 5) from a local politician in 2012 or 2013; 0 otherwise. Specifications employ pooled data to examine requests in either year. Cisterns treatment is coded 1 if respondent's household is in a neighborhood cluster selected for treatment; 0 otherwise. Rainfall shock is measured as the difference between rainfall in January-September of the relevant year and its historical municipal mean during identical months in 1986- 2011, divided by the municipality's historical monthly standard deviation of rainfall. Standard errors are clustered at the neighborhood level and reported in parentheses.

Table 4: Citizen Requests and Heterogeneity by Clientelist Relationship

	Request Any Private Good			Request Any Private Good Excluding Water	Request and Receive Any Private Good
	(1)	(2)	(3)	(4)	(5)
β_1: Cisterns Treatment	-0.012		-0.013	-0.017	0.005
	(0.013)		(0.013)	(0.012)	(0.010)
β_2: Cisterns Treatment × Monthly	-0.151***		-0.151***	-0.111**	-0.081**
	(0.049)		(0.049)	(0.045)	(0.036)
β_3: Cisterns Treatment × At-least Weekly	-0.063		-0.062	-0.022	-0.061*
	(0.042)		(0.042)	(0.039)	(0.033)
β_4: Rainfall Shock		-0.020*	-0.021**	-0.013	-0.007
		(0.011)	(0.010)	(0.009)	(0.009)
β_5: Rainfall Shock × Monthly		-0.026	-0.026	-0.017	-0.032*
		(0.026)	(0.026)	(0.024)	(0.019)
β_6: Rainfall Shock × At-least Weekly		-0.008	-0.005	-0.001	-0.014
		(0.018)	(0.018)	(0.017)	(0.015)
β_7: Monthly	0.167***	0.092***	0.166***	0.129***	0.085***
	(0.037)	(0.026)	(0.037)	(0.033)	(0.026)
β_8: At-least Weekly	0.088***	0.058***	0.088***	0.049*	0.069***
	(0.032)	(0.022)	(0.032)	(0.029)	(0.025)
Effect of Cisterns Treatment for Individuals in Clientelist Relationship:					
Monthly: $\beta_1 + \beta_2$	-0.163***		-0.163***	-0.128***	-0.076**
	(0.047)		(0.047)	(0.044)	(0.035)
At-least Weekly: $\beta_1 + \beta_3$	-0.075*		-0.074*	-0.039	-0.055*
	(0.041)		(0.041)	(0.038)	(0.032)
Effect of Positive 1 SD Rainfall Shock for Individuals in Clientelist Relationship:					
Monthly: $\beta_4 + \beta_5$		-0.046*	-0.047*	-0.031	-0.040**
		(0.026)	(0.026)	(0.024)	(0.020)
At-least Weekly: $\beta_4 + \beta_6$		-0.029	-0.026	-0.014	-0.022
		(0.018)	(0.018)	(0.016)	(0.014)
Test of Joint Hypothesis [p-value]:					
(a) $\beta_2 = \beta_3$	0.142		0.141	0.093	0.656
Municipality Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	4,288	4,288	4,288	4,288	4,284

Notes: Outcome variable is coded 1 if respondent reported: requesting a private good (columns 1- 3)/requesting a private good excluding water (column 4)/requesting and receiving a private good (column 5) from a local politician in 2012 or 2013; 0 otherwise. Specifications employ pooled data to examine requests in either year. Cisterns treatment is coded 1 if respondent's household is in a neighborhood cluster selected for treatment; 0 otherwise. Rainfall shock is measured as the difference between rainfall in January-September of the relevant year and its historical municipal mean during identical months in 1986- 2011, divided by the municipality's historical monthly standard deviation of rainfall. Standard errors are clustered at the neighborhood level and reported in parentheses.