Sensitivity Bias in Regime Support: Evidence from Panel Surveys in an Autocracy at War*

Timothy Frye†
& Henry Hale‡
& Ora John Reuter§
& Bryn Rosenfeld¶

September 3, 2024

Abstract

This paper reports findings from the first comprehensive study of sensitivity bias in regime support, and support for a cornerstone regime policy, in an autocratic setting. Our focus is support for Russian president Vladimir Putin and his war in Ukraine. We investigate potential sources of bias in survey estimates of crucial political attitudes that may arise from survey nonresponse, item nonresponse, and misreporting — making this the most wide-ranging study on the validity of survey-based measures of political support in Russia to date. With a unique three-wave panel study spanning Russia's full-scale invasion of Ukraine, we analyze patterns of attrition after February 2022, nonresponse to potentially sensitive questions, and — using three of the most common indirect methods of eliciting truthful responses to sensitive questions in the large literature on socially and politically stigmatized behaviors — whether Putin and the war in Ukraine are as widely supported as opinion polls suggest.

*Authors listed in alphabetical order. All authors participated in the design and fielding of the survey. Rosenfeld took the lead on designing the endorsement experiments and randomized response questions. Frye and Reuter took the lead on designing the list experiments. As of this version, Rosenfeld has taken the lead on the analysis and writing of the manuscript. A previous version of this paper was presented at the 2023 annual convention of the Association for Slavic, East European, & Eurasian Studies (ASEEES). For excellent research assistance, we wish to thank Francis Cayton, Davey Seeman, and Georgy Tarasenko. Frances contributed to all aspects of the analysis of preference falsification; Davey contributed to the analysis of panel attrition and item nonresponse; Georgy contributed to the analysis of the randomized response question.

†Marshall D. Shulman Professsor, Department of Political Science, Columbia University, Email: tmf2@columbia.edu. URL: https://www.timothyfrye.com/

‡Professor, Department of Political Science and International Affairs, George Washington University. Email: hhale@gwu.edu. URL: https://www.henryhale.net/

§Associate Professor, Department of Political Science, University of Wisconsin-Milwaukee, Email: reutero@uwm.edu. URL: https://www.ojreuter.com

¶Assistant Professor, Department of Government, Cornell University. Email: brr59@cornell.edu. URL: https://www.brynrosenfeld.com/

1 Introduction

Since Russia's full-scale invasion of Ukraine in February 2022, the Kremlin has enacted a range of measures to repress dissent and stepped up its repression of peaceful protest. The climate for publicly expressing political opinion, and thus for survey research in Russia, has changed. Yet Russian authorities have not sought to directly regulate household surveys. The Russian government has not pursued the Chinese model of controlling what topics or questions can be asked. Paradoxically, the war has brought more rather than fewer survey research entities into view in Russia: the government, independent, and opposition groups all conduct public opinion surveys in Russia today and publish their findings.

With these ongoing opportunities to field surveys in Russia come critical questions about the practice of polling in repressive environments. In this paper, we draw on new evidence from a probability-based nationally representative panel survey of Russian adults, conducted face-to-face in three waves that span Russia's full-scale invasion of Ukraine.[1]

2 Survey nonresponse and panel attrition

For many household surveys around the world, response rates have been steadily declining. Monetary incentives and multiple contacts are often used to increase survey response. Such costly strategies are used widely, because nonresponse poses such a serious threat to the accuracy of survey results. What is nonresponse? As Tourangeau & Plewes (2013) observe:

> Although there is relatively little ambiguity about what constitutes response, nonresponse may cover a broad range of possibilities beyond "refusal" and "no contact." Some nonresponse may reflect refusals that are so adamant that conversion is a virtual impossibility, but in other cases there is a degree of judgment about the utility of following up, perhaps with another interviewer or a specialized refusal convertor in an interviewer-mediated survey. Respondents who are not contacted might not be reached because they are away for

[1] More information about the survey can be found in Appendix A1.

a prolonged period, because contact was poorly timed, or simply because contact is random and more attempts would have been needed.

How has the war impacted potential respondents' willingness to be surveyed? The three-wave Russian Election Studies (RES) panel offers new evidence on this question. First, because the panel consists of two pre-invasion waves, we can examine whether the rate of attrition rose after the invasion,[2] as we might expect if potential respondents became more fearful. In fact, of the 2,501 respondents to the first wave of the survey (AAPOR RR-1 of 46%), conducted just before Russia's 2021 parliamentary election, 1,772 were reinterviewed approximately three months later, for a reinterview rate of 70.9 percent. This is on par with other election panel surveys in developing country contexts that pay respondents an incentive (e.g. the Argentine Election Study), as the RES did in each of its three waves.[3] In wave 3, approximately two years later, 1,038 respondents were reinterviewed, a reinterview rate of 58.6 percent. Additionally, of those who were eligible, interviewers were unable to recontact 248, making the completion rate among those who were successfully recontacted 68 percent (AAPOR COOP-1). Fully 82 percent of those interviewed in the latest wave after the invasion indicated – at the end of the survey after answering a wide variety of political questions – that they would be willing to participate in a future survey. In sum, then, we find no evidence that willingness to participate has fallen off sharply since Russia's invasion of Ukraine and subsequent tightening of political control.

Second, in a cross-sectional survey little or nothing is known about nonrespondents. Analyzing attrition from the panel, we have much more information about the political attitudes and characteristics of people who became nonrespondents in subsequent waves. In Table 1, we analyze the correlates of attrition in wave three of the panel, after the invasion of Ukraine. For ease of interpretation, all of the models are OLS. The first column examines differences in the demographic characteristics of attriters and nonattriters (i.e. those who drop out of the survey between waves two and three – after the invasion – and those who do not). The results suggest that men are more likely to drop

[2]We take as our measure of survey nonresponse panel attrition, though clearly there are differences in the attempt to initially recruit respondents and the attempt to reinterview them.)

[3]The incentive in the initial wave was approximately $5, rising to $15 in subsequent waves.

Table 1: Factors Predicting Panel Attrition

	Demographics (1)	Interest/ Engagement (2)	Regime Support (3)	Vote (prospective) (4)	Vote (2018) (5)	Risk attitudes (6)
Male	0.054**					
	(0.025)					
Age (logged)	-0.006					
	(0.033)					
Education	-0.094**					
	(0.042)					
Afford	0.049					
	(0.062)					
Lives alone	0.064*					
	(0.033)					
Town size	-0.016					
	(0.032)					
Employed	-0.029					
	(0.029)					
State-sector employment	-0.008					
	(0.033)					
Political interest		0.017				
		(0.037)				
Interpersonal trust		0.006				
		(0.042)				
Understanding of questions		-0.069				
		(0.077)				
Respondent competence		-0.059				
		(0.067)				
Attitude toward interview		-0.068				
		(0.080)				
Comfort during interview		0.028				
		(0.059)				
Putin disapproval			-0.021			
			(0.026)			
Putin approval DK			-0.062			
			(0.045)			
Opposition vote (prospective)				-0.092***		
				(0.028)		
Nonvoter (prospective)				0.039		
				(0.031)		
Undecided voter (prospective)				-0.005		
				(0.044)		
Opposition vote (2018)					-0.054	
					(0.036)	
Nonvoter (2018)					0.036	
					(0.027)	
Can't recall vote (2018)					0.0002	
					(0.044)	
Willingness to take political risks (wv 1)						0.028
						(0.049)
Willingness to take political risks (wv 2)						0.031
						(0.052)
Constant	0.464***	0.545***	0.425***	0.433***	0.409***	0.405***
	(0.137)	(0.080)	(0.015)	(0.018)	(0.018)	(0.016)
R^2	0.010	0.003	0.001	0.010	0.003	0.0006
Observations	1,744	1,741	1,772	1,772	1,772	1,661

*** p<0.01, ** p<0.05, * p<0.1 The table reports coefficients from OLS regressions predicting attrition in the third (post-invasion) survey wave. The reference group in column 3 is approval of President Putin (wave 2). In columns 4, it is indicating that one would vote for President Putin if an election were held on Sunday, while in column 6 it is having voted for President Putin in 2018. Source: RES surveys September and December 2021 and October 2023. The sample is all potential respondents to wave 3.

out of the sample as are the less educated and those who live alone. None of these differences are very large, and the patterns are consistent with evidence from panel studies in other contexts, including the U.S. (see e.g., Frankel and Hillygus, 2014).

Often, analyses of panel attrition focus only on demographic factors. We further investigate whether nondemographic factors might also be related to panel attrition. Following Frankel and Hillygus (2014), we first look at various aspects of survey experience and respondent disposition in the preceding panel wave. The model in column two includes standard measures of political interest and interpersonal trust[4] as well as four different interviewer evaluations of the respondent's engagement with the survey. The first asks the interviewer's perception of how well the respondent understood the questions, while the second inquires about their overall competence in taking the survey (e.g. whether they needed questions repeated). The third asks about the respondent's attitude toward the interview: whether they were friendly and interested, not especially interested, impatient and restless, or unfriendly. The fourth records the interviewer's impressions of the respondent's behavior during the interview: whether they were generally nervous, occasionally nervous, or comfortable. None of these factors are significant predictors of attrition and each of the estimated coefficients is close to zero. We thus find no evidence that low interpersonal trust or the interviewer's evaluation of a respondent's disposition are predictive of attrition (though most of the latter variables have the expected sign). It is worth noting that respondents were paid an incentive to participate in each wave of this study. This may have boosted participation among those with low political interest—as it was intended to do.

The next three columns of Table 1 explore whether political beliefs are related to survey participation in the post-invasion wave. Column 3 shows that whether a respondent approved of Putin before the full-scale invasion of Ukraine is unrelated to their willingness to be interviewed in wave 3. In column 4, we regress attrition in the post-invasion survey on a respondent's prospective presidential vote choice in 2021. The reference category is voting for President Putin. The coefficient on opposition voting implies that

[4]The question asked to what extent respondents agreed or disagreed with the statement: "Most people can be trusted."

respondents who said they would vote for an opposition candidate we 9.2 percentage points less likely to drop out of the panel than respondents who said they would vote for Putin. The pattern is similar in column 5, though the coefficient on voting for an opposition candidate in Russia's most recent presidential election in 2018 just misses conventional levels of statistical significance (p=.14). Contrary to the notion that Putin opponents are now too fearful to participate in surveys, leaving the views of Putin supporters to be disproportionately represented, these results imply that Putin's opponents are keen to have their voices heard. If anything, those who voted against Putin in the past and those who said in 2021 on the eve of the war that they would vote against Putin in the future were *less* likely to drop out of the sample.

Finally, column 6 examines whether respondents' risk attitudes are associated with their survey participation, as one would expect if survey participation is seen as a risky choice. For this analysis, we use a question in which respondents were asked in each preceding wave to describe their willingness to take risks when it comes to politics. We find no evidence that attitudes toward taking political risks is predictive of attrition. And, in fact, willingness to take risks in other life domains, like health and finances, was also unrelated to survey participation. Together these results cast doubt on the notion that opponents of the regime are increasingly unwilling to participate in surveys and that their nonparticipation is driven by the perceived risk of taking part. In fact, those who were already opponents of the regime were more, not less, willing to be reinterviewed as Russia's political climate became more repressive.

That said, this evidence does not preclude the possibility that some Putin voters became Putin opponents following the full-scale invasion, and that this group of respondents declined to be reinterviewed due to fear. This interpretation would require, however, that their attitudes toward risk were no different than others'. Note too that it would imply that prospective Putin voters who defected were fearful of being reinterviewed, while those who approved of Putin's performance in office but subsequently defected were not—since while Putin voting predicts attrition Putin approval does not. Finally, it would also require, for example, that they were not employed by the state, since state employees were no more likely than others to drop out. Given that employ-

ees of the state are more vulnerable than others to retaliation by authorities, it seems strange that they are not more likely to attrite—if fear is an important driver of survey nonresponse.

We can probe these interpretations further by comparing the results above with similar analysis of attrition in the initial pre-invasion waves of the study. If the factors that predict attrition between waves 1 and 2 are the same as those that predict attrition between waves 2 and 3, it would suggest that response patterns did not change appreciably after the full-scale invasion of Ukraine. By contrast, if a rising climate of repression affected who was willing to participate, we would expect to see differing patterns of attrition before and after February 2022. Table 2 replicates the attrition analysis presented in Table 1 for the second panel wave, in December 2021. The results are very similar with just a few exceptions.

First, whereas age was unrelated to attrition in the post-war survey, older respondents were more like to drop out of the panel between waves 1 and 2. Second, whereas none of the measures of respondent engagement were significant predictors of attrition in the post-invasion survey, two of the items coded by interviewers were associated with attrition prior to the war. Both how well the respondent understood the questions and the respondent's attitude toward the interview were significant predictors of attrition following the first wave of the panel—and in the expected direction. Thus, interestingly, the interviewer's impression that a respondent was impatient, anxious, or hostile was no better at predicting a respondent's subsequent willingness to be interviewed after Russia's political climate arguably deteriorated than it was before. With that said, the magnitude of the coefficients and their signs are quite similar before and after the invasion, though in the smaller post-invasion sample they lose statistical significance. Substantively, what this result suggests is that moving from the top to the bottom of the scale—that is, from "friendly and interested" to "hostile"—increased the likelihood of attrition by around 6 percentage points. This is a meaningful (if small) marginal effect, which could represent sample bias due to sensitivity.

Moving on, columns 3-5 show very similar patterns to those observed in the post-invasion survey. Neither regime support, measured as Putin approval, nor a respon-

Table 2: Comparison of Factors Predicting Panel Attrition Pre-Invasion

	Demographics (1)	Interest/ Engagement (2)	Regime Support (3)	Vote (prospective) (4)	Risk attitudes (5)
Male	0.064***				
	(0.019)				
Age (logged)	0.085***				
	(0.026)				
Education	-0.028***				
	(0.008)				
Afford	0.062				
	(0.046)				
Lives alone	0.055**				
	(0.025)				
Town size	-0.005				
	(0.006)				
Employed	-0.022				
	(0.022)				
State-sector employment	0.032				
	(0.025)				
Political interest		0.014			
		(0.009)			
Interpersonal trust		0.002			
		(0.011)			
Understanding of questions		-0.081***			
		(0.027)			
Respondent competence		-0.002			
		(0.017)			
Attitude toward interview		-0.057***			
		(0.019)			
Comfort during interview		-0.027			
		(0.022)			
Putin disapproval			-0.025		
			(0.020)		
Putin approval DK			-0.048		
			(0.035)		
Opposition vote (prospective)				-0.036*	
				(0.021)	
Nonvoter (prospective)				-0.069***	
				(0.025)	
Undecided voter (prospective)				-0.016	
				(0.039)	
Willingness to take political risks (wv 1)					-0.011
					(0.034)
Constant	0.019	0.769***	0.303***	0.317***	0.293***
	(0.107)	(0.085)	(0.012)	(0.014)	(0.012)
R^2	0.020	0.018	0.001	0.003	3.99×10^{-5}
Observations	2,452	2,457	2,501	2,501	2,428

*** p<0.01, ** p<0.05, * p<0.1 The table reports coefficients from OLS regressions predicting attrition in the second survey wave (prior to the invasion). The reference group in column 3 is approval of President Putin (wave 1). In columns 4, it is indicating that one would vote for President Putin if an election were held on Sunday. Information on past presidential vote choice in 2018 is not available in wave 1, so the equivalent of column 5 of Table 1 is omitted. Source: RES surveys September and December 2021.

dent's self-assessed willingness to take political risks were systematically associated with survey participation. Meanwhile, opposition voters, the results in column 4 show, were already less likely to drop out of the panel before the war than were Putin voters, consistent with the interpretation (as above) that such respondents derive greater expressive value from survey participation. Relative to those who intend to vote for a candidate other than Putin and those who do not intend to cast a ballot,[5] Putin voters appear as either the lazy or hesitant survey takers.

Of course, the patterns of survey nonresponse may be different in different modes of interview (e.g., phone, online), in non-panel surveys in which interviewers have not built rapport with respondents, or in polls where respondents are not compensated for their participation. However, having detailed information on respondents' characteristics, attitudes, and past political behaviors from prior survey waves, as we do in this panel, gives us unique insight into the choice to participate in political surveys in a repressive environment. No other survey in Russia since the start of the war in Ukraine offers such evidence. Tentatively given the preliminary nature of this analysis, the evidence here does not suggest that regime opponents are reticent about survey participation. Though other interpretations are possible, one has to tell a very convoluted story.

[5]The result for non-voters differs before and after the invasion. Whereas respondents who indicated that they would not vote in a presidential election in 2021 were significantly less likely than Putin voters to attrite from the panel, in 2023 the relationship was not significant, and the sign flipped.

3 Item nonresponse

Besides survey nonresponse, nonresponse to individual survey questions (or "item nonresponse") may introduce bias in the estimation of outcomes of interest. Research by Shen and Truex (2021) shows that while, in many authoritarian countries, citizens are about as likely to avoid questions about their government, democracy, and respect for human rights as citizens in democracies, self-censorship is higher in countries with the most closed political systems. This raises the possibility that nonresponse may have grown as Russia's political climate has deteriorated, concealing opposition to the regime and its war.

Reisinger, Zaloznaya, and Woo (2023) discuss the issue of question-specific (item) nonresponses and how they introduce biases that undermine inference, drawing on survey data from Russia and neighboring countries prior to the invasion. Based on an analysis of sensitive questions regarding corruption involvement, they conclude that nonresponse and misreporting (which we deal with below) were already a problem before Russia's more repressive turn during the Ukraine war. We share their interest in identifying strategies to reduce item nonresponse and addressing it when interpreting survey data. We turn to the evidence from our own survey next.

3.1 Have "Don't know" responses risen? How much?

We begin the analysis by asking whether "don't know" responses and refusals to answer questions about regime support have risen and, if so, by how much. We focus on regime support as the most indicative example of a potentially sensitive question. Note that with only one post-invasion survey wave, we leave aside the question of whether nonresponse has risen on the question of war support over the nearly two years since Russia's full-scale invasion of Ukraine. Figure 1 shows the pattern of nonresponse to a binary question about approval of President Putin before and after the war. While there was no difference in the nonresponse rate between the first two surveys waves (again, the surveys were conducted only 3 months apart at the end of 2021), there is a sharp rise in nonresponse in the third survey wave after the full-scale invasion of Ukraine. This

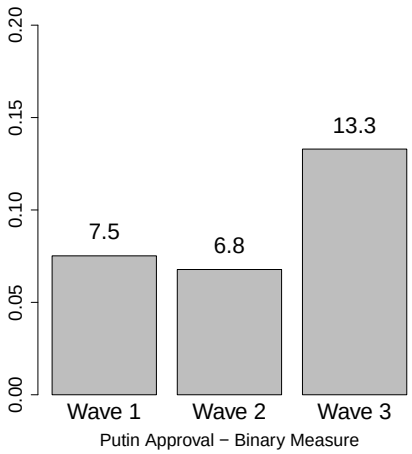

Figure 1: Nonresponse to Putin approval question has risen.
This figure displays the rate of nonresponse to a binary question about approval for President Putin (including "don't know" and refusals to respond) across the three survey waves. The difference between the first and second waves is not significant. However, there is a meaningful and statistically significant (p<.001) increase in nonresponse to 13.3 percent in the third survey wave after the full-scale invasion of Ukraine.

increase of about 6-7 percentage points is statistically significant at p<.001.

Consider a second piece of evidence from the same survey, pictured in Figure 2. In addition to asking respondents whether they do or do not approve of Putin's performance in office, a binary formulation, we also asked respondents to give their views of Putin on a 0-10 scale, from "entirely unfavorable" to "entirely favorable." Here, while there is again a statistically significant (p<.01) increase in the rate of nonresponse in the third survey wave, substantively the increase is small and the overall level of nonresponse is much lower, just 3.5 percent. The extent of nonresponse thus varies by question type, across questions designed to probe the same potentially sensitive underlying concept: support for President Putin. Whereas the potential level of evasive nonresponse—in other words, effort by respondents to avoid saying what they really think for fear of reprisal—could be greater than 10 percent according to the binary measure, it is just a few percentage points according to the thermometer.

3.2 Measuring Uncertainty: It's the Questions, Stupid!

Whether, however, the rise in "don't knows" just documented represents evasive nonresponding is another question. Rather than fear driving nonresponse, respondents' "don't

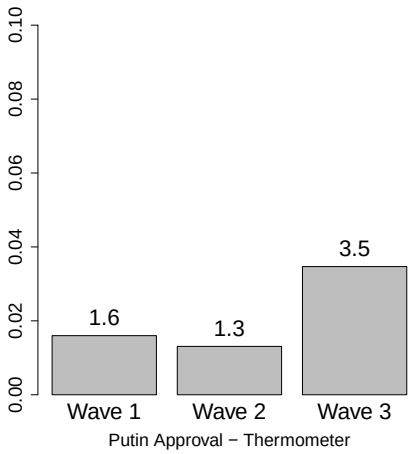

Figure 2: Nonresponse to Putin feeling thermometer has risen, hardly at all.
This figure displays the rate of nonresponse to the Putin feeling thermometer (including "don't know" and refusals to respond) across the three survey waves. The difference between the first and second waves is not significant. While there is a statistically significant (p<.01) increase in the rate of nonresponse in the third survey wave after the full-scale invasion of Ukraine, substantively the increase is small and nonresponse remains low.

knows" could mean just that: that their views are uncertain. One appraoch to answering this question is to explore the correlates of nonresponse, a task we turn to next.

Figure 3 and Figure 4 compare the full distribution of responses on dichotomous versus thermometer measures of Putin approval and war support, respectively. Nonresponse is fully 10 percentage points lower on both questions when they are asked using the more extensive response scale. This raises the possibility that the dichotomous response scale cannot accommodate a sizable share of respondents whose views on both subjects are perhaps not sensitive, but equivocal. In other words, the problem could be the questions rather than the respondents.

For the Putin approval question, we see that when given the option, fully 11 percent of respondents locate themselves at the mid-point of the scale. On the item about support for the 'special military operation,' 13 percent select the mid-point of the scale. One interpretation, then, is that to the extent that there is rising nonresponse, it reflects growing uncertainty in respondents' views, rather than fear. This conclusion is echoed by Nadia Evangelian and Andrei Tkachenko, who argue that "don't know" responses more likely reflect respondents' lack of clear opinions on the war (and other political issues) than fear of expressing opposition, based on their analysis of data from six post-war

waves of the Chronicle survey.[6]

Figure 3: Nonresponse to Putin Approval Question Varies by Question Type

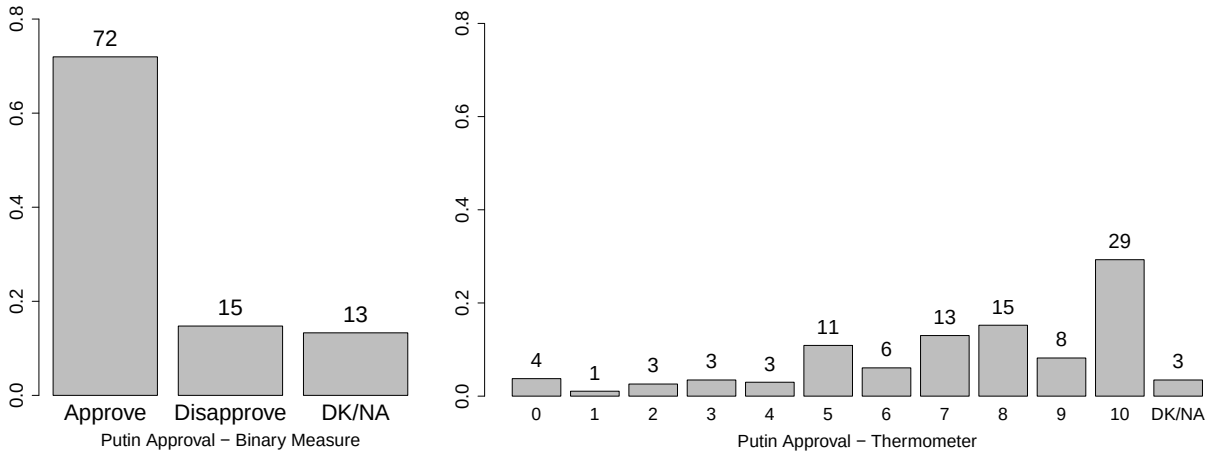

Note: The left panel displays responses to the question on approval of President Putin asked with binary response options ("Approve," "Do not approve") and the rate of nonresponse. The right panel displays responses to the 0-10 thermometer rating of President Putin. The rate of nonresponse is significantly lower (10 percentage points). Fully 11 percent of respondents locate themselves at the mid-point of the scale.

Of course, it is possible that, when given the chance, respondents hide their opposition in the mid-point of the scale, rather than choose to say "don't know" or refusing to answer the question. In this case, respondents would be answering strategically, misrepresenting their preferences, a subject we turn to next. In fact, there is some evidence for this in Table 3.

Table 3 regresses nonresponse to both binary and thermometer versions of the Putin approval and war support questions on respondents' risk attitudes. If risk-averse respondents are more likely to answer sensitive political questions evasively by choosing "don't know" or refusing to answer the question, we would expect to observe a significant negative coefficient on risk attitudes. There is absolutely no evidence of this for the Putin support questions, as can be seen in columns 1 and 2. On the contrary, people who said they were more willing to take political risks in the pre-invasion survey (wave 2) were more likely to answer "don't know" when asked whether they approved

[6] https://www.extremescan.eu/post/14-the-first-phase-of-a-special-military-operation-in-the-minds-of-russians

Figure 4: Nonresponse to War Support Question Varies by Question Type

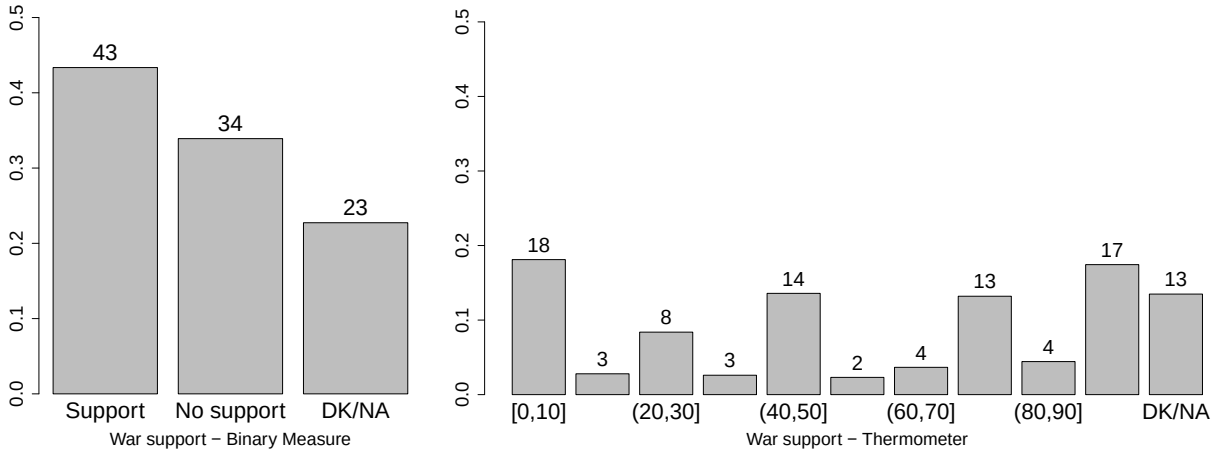

Note: The left panel displays responses to the question on support for the war in Ukraine asked with binary response options ("Support," "Do not support") and the rate of nonresponse. The right panel displays responses to the 0-100 thermometer rating of support for the war. The rate of nonresponse is significantly lower (again, by 10 percentage points). By this measure, 13 percent of all respondents locate themselves at the mid-point of the scale.

of Putin's performance in office. The evidence again implies that their nonresponse was more likely due to increasing uncertainty rather than the risks of expressing opposition.

In the war support questions, however, the pattern is different. As shown in column 3, there is a negative and statistically significant association between risk attitudes measured in the survey's first wave and contemporaneously in the third wave and nonresponse to the dichotomous war support question. This implies that respondents who were more averse to taking political risks were more likely to be nonresponders to the war support question, when they lacked other options (given the dichotomous formulation). Column 4 then shows that they were not, however, more likely to be nonresponders to the war support thermometer. Two interpretations are possible: either risk aversion is correlated with more indecision (though note, only on support for the war and not support for Putin) or risk averse respondents choose "don't know" when they have no place to hide and respond strategically at the middle when given a more extensive scale. In fact, regressing risk attitudes on the war support question, we find that being more risk averse is associated with responding at the middle of the war support scale, though the result is just shy of statistical significance ($p = .14$). Again it is plau-

sible that risk aversion is associated with more equivocal views on the war; though it is difficult to explain why it would be associated with more equivocal views on the war but not Putin. Alternately, risk averse respondents may indicate "don't know" to a binary war support question but report uncertainty/indifference on a thermometer of war support as a safer alternative to expressing discontent.

Table 3: Relationship Between Risk Attitudes and Item Nonresponse

	Putin Approval (binary) DK/NA (1)	Putin Thermometer DK/NA (2)	War Support (binary) DK/NA (3)	War Support Thermometer DK/NA (4)	War Support Thermometer Middle response (5)
Risk Aversion (wv 1)	-0.010	0.035	0.107*	0.031	0.096*
	(0.046)	(0.021)	(0.059)	(0.047)	(0.056)
Risk Aversion (wv 2)	-0.094*	-0.012	-0.013	0.051	-0.047
	(0.049)	(0.022)	(0.062)	(0.050)	(0.059)
Risk Aversion (wv 3)	0.031	0.002	0.100*	0.058	0.049
	(0.044)	(0.020)	(0.056)	(0.045)	(0.053)
Constant	0.104***	0.028***	0.252***	0.150***	0.205***
	(0.016)	(0.007)	(0.020)	(0.016)	(0.019)
R^2	0.005	0.003	0.010	0.006	0.005
Observations	895	895	895	895	895

*** p<0.01, ** p<0.05, * p<0.1 The table reports coefficients from OLS regressions on nonresponse to the regime and war support items in the third (post-invasion) survey wave. All explanatory variables have been rescaled between 0 and 1. Source: RES surveys September and December 2021 and October 2023. The sample is all wave 3 respondents.

3.3 Correlates of Non-Response

While this evidence suggests that uncertainty rather than fear accounts for much of the item nonresponse in our survey, even in a more repressive political environment, we can look more broadly, beyond these indicative examples of potentially sensitive regime questions. For a preliminary analysis, we coded as either "sensitive" or "nonsensitive" each of the questions in our survey.[7] The total number of questions was 121. We coded 82 of these as potentially sensitive and 39 as nonsensitive.

Figure 5 plots the distribution of nonresponse for sensitive and nonsensitive items, based on the full universe of questions in our post-invasion survey. The y-axis gives

[7]The first step in doing so was to query ChatGPT. Subsequently, a skilled undergraduate research assistant revised the lists of sensitive and nonsensitive items produced by ChatGPT. The next step in this analysis is to draw up a coding rubric and have multiple research assistants manually code these items, ensuring a high level of intercoder reliability.

the total number of nonresponses for each question (i.e. the nonresponse count) out of a potential 1,038. The interquartile range and median are also plotted. While the highest recorded level of nonresponse (425) was to a nonsensitive question ("Do you or do you not approve of the activities of Xi Jinping?"),[8] the median level of nonresponse was somewhat higher on potentially sensitive questions than on nonsensitive questions (82 vs. 33, or approximately 8% vs. 3%). While reticence could explain this higher nonresponse rate for potentially sensitive questions, it could also reflect the relatively greater complexity of the sensitive items, the majority of which are attitudinal questions, compared with the mostly demographic nonsensitive items.

To probe these interpretations further, we next examine the drivers of nonresponse for potentially sensitive and nonsensitive question using regression analysis. If the factors associated with nonresponse to the potentially sensitive items are the same as those associated with nonresponse to the nonsensitive items, it would imply that something other than fear is driving observed nonresponse patterns. Table 4 presents results from OLS models in which we regress respondent characteristics on nonresponse. The dependent variable in the first column is the count of nonresponses to nonsensitive items. In the second column, it is the count of nonresponses to the potentially sensitive items. Both dependent variables have been rescaled 0-100 for comparability.

Overall, the characteristics associated with nonresponse are very similar for the potentially sensitive and nonsensitive items. Moreover, in several cases, the observed patterns of nonresponse are actually inconsistent with the argument that nonresponse is driven by fear. Turning to the results, we find several patterns familiar from the nonresponse literature and surveys in less repressive, democratic contexts. Women have higher rates of nonresponse to both nonsensitive and sensitive items, as do respondents who are less politically interested and respondents who are lower in interpersonal trust.[9]

[8]Sensitivity bias can arise from either political or social pressure to conform, i.e. from the fear of sanction by ones peers for deviating from an accepted norm or from fear of sanction by the state. We reason that it is unlikely respondents would expect their peers to care about their views on the Chinese premier, just as it is unlikely that they would expect the state to retaliate for unfavorably views of the Chinese leader—even if China is Russia's political ally.

[9]A significant negative coefficient on interpersonal trust for sensitive item but not nonsensitive items would be more concerning. The similarity of the coefficients, however, suggests a pattern of responding that has little to do with the putative sensitivity of a question.

Table 4: Factors Predicting Item Nonresponse

	Nonsensitive Items	Sensitive Items
Male	-0.985^{*}	-1.384^{*}
	(0.421)	(0.583)
Age 18-29	1.359	1.495
	(0.898)	(1.171)
Age 30-39	1.291	1.513
	(0.743)	(1.153)
Age 40-49	0.570	0.962
	(0.755)	(1.125)
Age 50-59	0.110	0.129
	(0.594)	(0.937)
Less than secondary edu	-1.162	-1.736
	(0.668)	(1.137)
Complete secondary	0.488	0.561
	(0.717)	(0.935)
Specialized secondary	0.944^{*}	0.366
	(0.423)	(0.728)
Incomplete higher	-0.524	-0.795
	(0.924)	(1.592)
Afford	-0.914	-1.056
	(1.171)	(1.603)
Lives alone	-0.709	-0.699
	(0.490)	(0.823)
Moscow/Major City	-1.154	0.144
	(0.725)	(1.006)
City of 100k-500k	-2.340^{*}	-3.320^{**}
	(0.952)	(1.202)
Urban settlement $<$ 100k	-2.425^{***}	-2.833^{**}
	(0.608)	(0.956)
Employed	-0.160	-0.194
	(0.699)	(0.936)
State-sector employment	-1.122^{*}	-1.340
	(0.517)	(0.837)
Political interest	-1.714^{**}	-5.650^{***}
	(0.633)	(0.891)
Interpersonal trust	-3.247^{***}	-2.854^{*}
	(0.819)	(1.211)
Attitude toward interview	-0.189	-0.841
	(0.973)	(1.463)
Understanding of questions	-0.661	-1.590
	(0.694)	(1.258)
Comfort during interview	2.147^{*}	1.971
	(0.858)	(1.694)
Respondent competence	0.225	0.278
	(0.693)	(1.371)
Putin disapproval	-1.001^{*}	-3.534^{***}
	(0.444)	(0.682)
R^2	0.228	0.306
Num. obs.	966	966

*** p<0.001, ** p<0.01, * p<0.05 The table reports coefficients from OLS regressions predicting item nonresponse in the third (post-invasion) survey wave. All models include region (oblast) fixed effects. Source: RES survey October 2023. The sample is all wave 3 respondents.

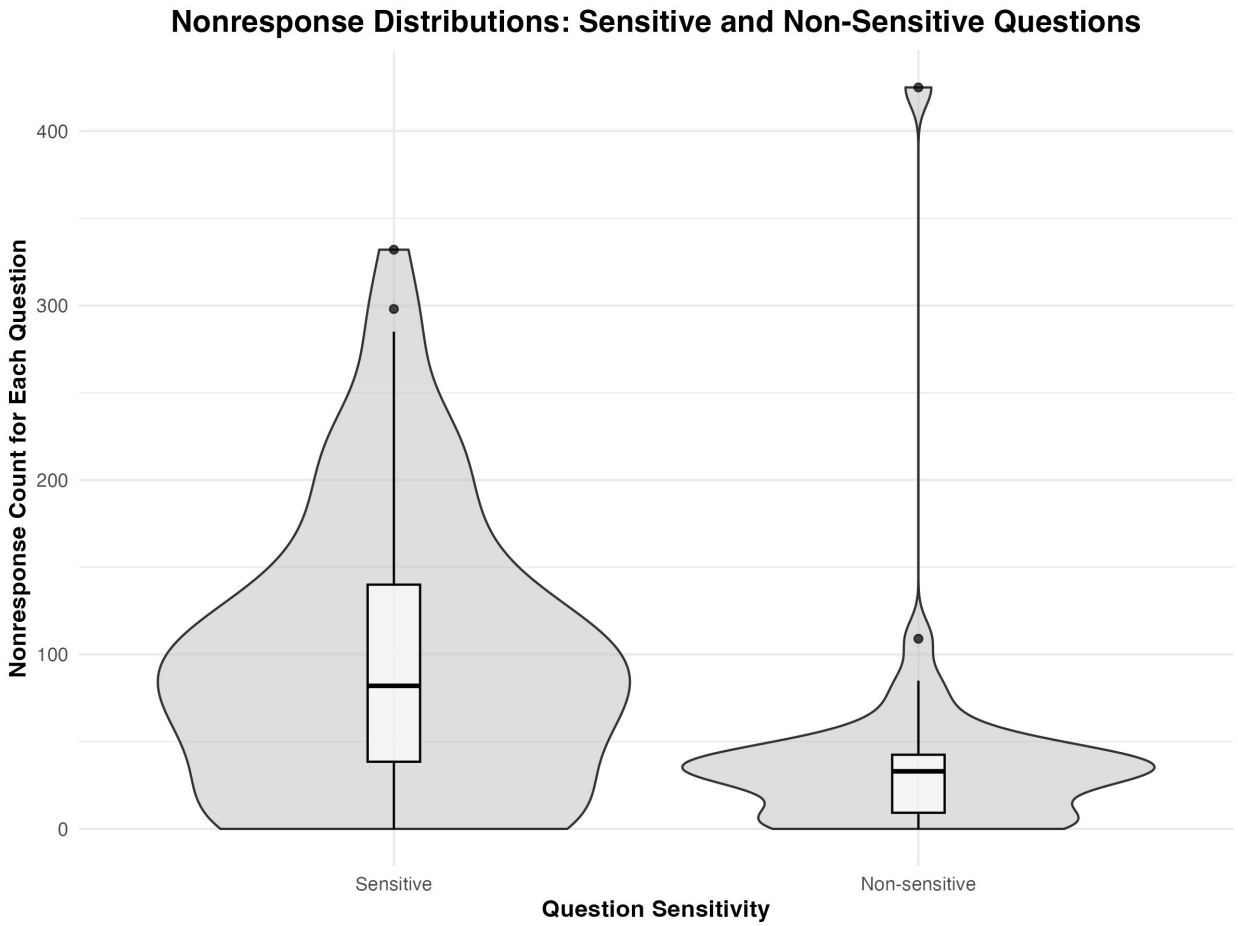

Figure 5: Distribution of item nonresponse by question sensitivity.
This figure displays, as a violin plot, the total number of nonresponses for each question, as well as the median and interquartile range, separately for sensitive and nonsensitive questions.

There is also some evidence that people in urban areas are less likely to decline to answer than are people living in villages.

Notably, state-sector employees who are more vulnerable to retaliation actually have *lower* levels of nonresponse (though the coefficient is similar in both models, it misses conventional levels of significance for potentially sensitive items). The same is true of those who express their disapproval of President Putin: they are *less* likely to decline to answer both questions that are potentially sensitive and those that are not (though the magnitude of the effect is somewhat larger for sensitive items). This could be because Putin opponents are particularly courageous or because fear of repercussions has little to do with response patterns, and those who oppose Putin also hold more definite views than other respondents on a variety of political questions. In sum, these patterns appear

more consistent with nonresponse that reflects uncertainty in respondents' views and is driven by political interest and other respondent characteristics (e.g., gender) rather than reticence due to question sensitivity.

The one counterintuitive result in Table 4 is for respondent comfort. As a reminder, this is an interviewer-coded assessment of the respondent's comportment during the interview: from nervous to comfortable. While significant only for nonsensitive items, the positive coefficient implies that greater comfort was associated with more nonresponse.[10] If nervous respondents were indeed more likely to answer survey questions (though note that there is no statistically significant evidence of this for sensitive items), then it would be especially important to gauge the sincerity of responses—as we do in the next section.

On the whole, then, much of this analysis points to real uncertainty of opinion as a diver of item nonresponse. Though this evidence is apparently at odds with Reisinger, Zaloznaya, and Woo's (2023) conclusion that fear of punishment is the key driver of nonresponse in Russia, there is at least one important difference worth bearing in mind. This concerns the legality of the behavior respondents are asked to report. Officials soliciting a bribe or respondents offering them to state agents are both illegal behaviors. Expressing disapproval of Putin or other officials is not, even if it is socially undesirable—and surely acknowledging that Putin has been vicious toward his most vocal opponents. Expressing opposition to the war is clearly a gray area, given the criminalizing of dissent (under the so-called "fake news" laws passed after the invasion) and many cases brought by Russian prosecutors thereafter. Still, saying that one does not support the 'special military operation' is not an explicitly criminal act like giving a bribe. And there are no known cases of a survey respondent being prosecuted for expressing opposition to the war in the context of a survey interview. Tentatively, much of the "problem" of nonresponse in Russian surveys appears from the perspective of this analysis to be a function of questions whose design does not allow an alternative to nonresponse when a respondent's views are uncertain or conflicted.

[10] Moving from the bottom to the top of the three-point scale increased the number of nonresponses by two.

4 Preference falsification

Still, we must consider further the issue of preference falsification, or the tendency of respondents to misrepresent their views due to political or social pressure (Buckley et al., 2023). Even if there is little bias from nonresponse – including survey nonresponse and item nonresponse – social desirability bias may distort estimates of key outcomes of interest, like regime support or support for regime policies like war. In other words, respondents who participate in surveys and answer interviews' questions may still dissemble about their true preferences. As political or social pressure to express a particular opinion grows, respondents may in fact become less likely to decline to answer pollsters' questions and more likely to misrepresent their views.

Most survey questions are asked directly, as in "Do you support or do you not support the special military operation in Ukraine?" or "Do you approve or disapprove of Vladimir Putin's activities as President of Russia?" And most of what the media report in Russia and abroad about Russians' support for Vladimir Putin and the war is based on direct questions. However, a large body of research shows that direct questions can lead to substantial underreporting on sensitive topics. To deal with these challenges, researchers across a several disciplines have developed alternative approaches to asking sensitive questions that provide respondents' additional privacy and encourage greater candor.

The most widely used of these methods in political science by far has been the list experiment (or Item Count Technique). Other indirect questioning techniques used to reduce social desirability bias in surveys include endorsement experiments and the randomized response technique. Evidence from validation studies that compare estimates obtained by asking questions indirectly, in ways that protect respondents by veiling their individual responses on the sensitive issue, suggests that indirect questioning techniques provide a picture that is closer to the truth (Rosenfeld, Imai, and Shapiro, 2016). Moreover, according to a recent meta-analysis which compared estimates from published and unpublished list experiments in political science to estimates from direct questioning, regime support in nondemocratic settings is one of the few topics widely studied in

political science[11] in which sensitivity bias is large enough that list experiments are preferred to direct questions (Blair, Coppock, & Moore 2020).

Russia's war in Ukraine has of course been ongoing since 2014. This longer period since the annexation of Crimea has witnessed a tightening of political control and deteriorating climate for free expression. Even before the full-scale assault on Ukraine began in February 2022, there has been a need for caution in interpreting direct questions about Putin and his policies. Mixed evidence on the sensitivity of political questions in Russian opinion polling has existed for some time. On the one hand, there is evidence that Russian survey respondents have been less fearful, less prone to lie than commonly assumed. Frye et al. (2017) determined that Putin's high approval ratings following Crimea's annexation were largely sincere. A recent paper by the same scholars interprets a raft of new evidence as suggesting that Putin's support remains largely genuine — though they also note, on the other hand, that "there is considerably more uncertainty [today] about Putin's true support than was apparent in 2015" (Frye et al., 2023). And, indeed, even the sincerity of support for Putin after Russia annexed Crimea has recently come into question. Hale (2022), based on new analysis of Russian surveys done in 2015, several months after Frye et al.'s (2023) study, finds that misrepresentation was an important factor in the surge in Putin's post-Crimea approval rating.

Research in social psychology suggests that the appearance of broad support begets even greater support as people cue off others or strive to fit in. The Kremlin's weaponization of polling since the Ukraine invasion exploits this fact, as Alyukov (2024) explains. Recent survey evidence from Russia suggests that surging support for the war may also have been in part be insincere. Chapkovski and Schaub (2022) find, using a list experiment, that support for Putin's 'special military operation' may have barely reached a majority, even in spring of 2022, and that direct questioning inflates support by approximately 10 percentage points. In this paper, we probe the stability of that finding, looking for evidence of social desirability bias in support for the war in Ukraine nearly two years after the full-scale invasion.

[11] Other topics included vote buying, voter turnout, racial prejudice, religious prejudice, and prejudicial views on sexual orientation.

We make several contributions. First, to the best of our knowledge, this is the first study to examine the issue of sensitivity bias in war support using a probability-based nationally representative survey. This is important since online samples which are younger and more educated than the population may overstate or understate sensitivity bias in overall support for the war. Second, in addition to conducting a list experiment on war support, we implement placebo tests, as recommended by Frye et al. (2023) to help diagnose the potential for artificial deflation, a design effect that can arise in list experiments. Understanding the potential for artificial deflation in our given research context is crucial, since artificial deflation may make support appear lower, and thus sensitivity bias higher, than it is in fact. Previous studies of Russians' war support that employ list experiments have not included placebo tests. Third, as far as we know this is the first study to investigate sensitivity bias in either support for the regime or its war that has done so using multiple indirect methods. Here we triangulate responses across the three most commonly used indirect questioning methods in the large literature on socially and politically stigmatized behaviors: the list experiment, the endorsement experiment, and the randomized response technique.

We next briefly explain how each technique works.

4.1 Endorsement Experiment

The endorsement experiment works by exploiting evaluation bias in human judgement. A sample of respondents is randomized into two groups. In the control group, respondents are asked to evaluate some relatively uncontroversial issue or object (e.g., rate a policy on a Likert scale). In the treatment group, that issue or object is associated with the sensitive item before being evaluated (e.g., the same policy is said to be endorsed by a controversial political leader). The difference between these two groups is then taken to reflect the degree to which respondents are favorable (or unfavorable) towards the sensitive item.

Here, we use the endorsement experiment to measure support for a leader (Putin) as well as a sensitive political behavior (supporting the war). To measure war support, a behavior, we flip the standard design, following Rosenfeld, Imai, and Shapiro (2016),

and ask respondents to rate their support for political actors (which is relatively uncontroversial) and then randomize the pairing of those actors with support for the sensitive outcome of interest. If this pairing induces a positive effect on support for the political actors, we interpret this effect as evidence that they support the war.

Specifically, regarding support for President Putin, we asked the control group:

```
Tell me, how much do you support the following initiatives?

A. The World Health Organization recommends that all adults should be
vaccinated against COVID-19. How do you feel about this proposal?

B. A recent proposal calls for expanding the system of medical check-ups
and periodic screenings, taking into account the current epidemiological
situation.

C. It has been proposed to use market mechanisms to ensure the
predictability of prices rather than setting prices from the top.

D. It has been proposed to create a carbon utilization sector, bring down
emissions and introduce strict control and monitoring measures.

    Completely support
    Mostly support
    Mostly do not support
    Completely do not support
    Hard to say
    Refuse to answer
```

In the treatment group, the statements read as follows, with the added information about President Putin's endorsement:

```
A. President Putin has endorsed the World Health Organization's
recommendation that all adults should be vaccinated against COVID-19.

B. A recent proposal by President Putin calls for expanding the
system of medical check-ups and periodic screenings, taking into
account the current epidemiological situation.

C. President Putin has proposed to use market mechanisms to ensure
the predictability of prices rather than setting prices from the top.

D. President Putin has proposed to create a carbon utilization sector,
bring down emissions and introduce strict control and monitoring measures.
```

Regarding support for the war, we asked the control group:

```
We would like to know your general opinion about some public figures.
I will mention a name, and please tell us your opinion about this
individual - very favorable, somewhat favorable, somewhat unfavorable,
or very unfavorable.

Prime Minister Mikhail Mishustin?
Minister of Foreign Affairs Sergey Lavrov?
Chairman of the State Duma Vyacheslav Volodin?

  Very favorable
  Somewhat favorable
  Somewhat unfavorable
  Very unfavorable
  Hard to say
  Refuse to answer
```

And in the treatment group, we added the information that each public figure supported the "special military operation,"

```
..., who expressed his full support for the SMO in Ukraine.
```

The endorsement experiment is grounded in extensive research on persuasion in social psychology (see Petty and Wegener, 1998, for a review). Researchers have found that individuals are more likely to be persuaded and influenced by likable sources (Petty and Cacioppo, 1986; Cialdini, 1993) and that endorsements of policies and positions are much more effective when an individual has positive affect toward the source of the endorsement (Wood and Kallgren, 1988; Chaiken, 1980; Petty, Cacioppo, and Schumann, 1983). As O'Keefe (1990) summarizes, "Liked sources should prove more persuasive than disliked sources" (p. 107).

The main advantage of the endorsement experiment is that, unlike the list experiment, it can never reveal the truthful answer to the sensitive question. However, this indirect nature also presents a major drawback in that a latent variable model is needed to derive estimates of sensitive behaviors from the ordered responses (as discussed in Bullock, Imai, and Shapiro 2011), and the endorsement effects do not have an obvious

scale.[12] The endorsement experiment is also statistically inefficient, even when compared with other indirect questioning techniques. For this reason, Bullock, Imai, and Shapiro (2011) recommends that the researcher use multiple questions to measure one sensitive item. The design we use here, with three to four items, still yields less statistically efficient estimates relative to other approaches.

4.2 List Experiment

Under the standard design of the list experiment, as in the endorsement experiment, researchers randomize a sample of respondents into two groups. A list of several control items is presented to the control group, and a list of the same control items plus one sensitive item of interest is read to the treatment group. Respondents are then asked to count the number of items on their list that fit certain criteria rather than respond to each item separately. The difference in means between the two groups then provides the simplest estimate of the prevalence of the sensitive attitudes or behavior in a target population though more efficient estimators are now available.

In our study, we used this standard design and asked the following question to the control group to measure Putin's support:

```
Now we'd like you to take a look at this list of names with political leaders.
Tell me how many of them you approve of, in general. Do not list specific
names - only say a number from 0 to 4 that indicates how many of these
leaders you approve of.

    1. Joseph Stalin
    2. Leonid Brezhnev
    3. Boris Yeltsin
    4. Peter I (the Great)
```

For the treatment group, the same exact script was read but the following additional sensitive item of interest was added to the list:

[12] They can, however, be benchmarked against the effect for endorsers whose level of support is commonly understood to be strongly positive or negative.

5. Vladimir Putin

To measure war support, we asked:

Look at this list of statements and say how many of the following things are true of you. Do not name specific statements - just tell me a number from 0 to 4.

1. My father did not have higher education.
2. I watch TV, Youtube or streaming services (IVI, OKKO, Kinopoisk, etc.) at least once a week.
3. I have visited Moscow.

For the treatment group, the same exact script was read but the following additional sensitive item of interest was added to the list:

4. I support the continuation of the special military operation.

Respondents in each group could report the number of items that applied to them and, like the direct question, they also had "don't know" and "no response" options.

The advantage of this approach is that respondents do not directly report whether the sensitive item applies to them. Instead, they provide a count of items on a list which contains other items. The major limitation of the list experiment is the problem of ceiling and floor effects. Answering "0" to our list experiment in the treatment group, for example, reveals that the person does not support Putin. To address these concerns, Glynn (2013) recommends that the the researcher choose control items such that responses to those items are negatively correlated. Other disadvantages are the fact that aggregation in the list often decreases the statistical efficiency of subsequent analyses and adding a sensitive item may alter one's (latent) response to control items (Blair & Imai 2012).

4.3 Randomized Response Technique

The randomized response method obscures individual responses by introducing random noise. A number of designs have been introduced since the work of Warner (1965)

(see Blair, Imai & Zhao 2015 for a review). In our study, we adapt the standard forced response design where a coin flip is used for randomization, but replace the coin flip with the birthday method.[13] Because the randomized response is thought to be difficult for respondents to grasp, we first gave respondents a chance to practice. We then proceeded to ask the main question about their opposition to the SVO (we did not ask a randomized response question about support for Putin). Our script is given here:

To answer the following question, first let's practice. Think what date your birthday falls on - is it even or odd. Now please answer "no" if you do not have higher education OR if your birthday is on an odd date. Otherwise (only if you have higher education AND your birthday is on an even date) answer "yes."

- Yes
- No
- Don't know
- Refused

Now please think about the date of your mother's birthday - is it even or odd. In this question, please do not tell me what day it is. Simply answer "no" if you do not support the special military operation in Ukraine OR if your mother's birthday is on an odd date. Otherwise (only if you support the SVO AND your mother's birthday falls on an even date) answer "yes."

- Yes
- No
- Don't know
- Refused

To discourage break-offs, interviewers were instructed to use the following script if respondents expressed confusion or hesitation about the answering the question:[14]

Sometimes, respondents would like to keep their answers private even from the interviewer. We will ask you a question in such a way that no one but you knows the answer to it. To do this, you only need to remember the date and understand whether it falls on even or odd number.

[13] Whether a respondent's birthday falls on an even or odd date is taken as quasi-random, with the probability of an odd birthday .49 and the probability of an even birthday .51.

[14] Interviewers were also instructed to clarify that if the respondent did not have a mother or could not recall her birth date to think of the birth date of another close relative or friend.

A main disadvantage of randomized response is the burden it imposes on respondents. Here, the method also requires respondents to administer randomization on their own, and this can lead to a high rate of refusal and attrition. Both Coutts and Jann (2011) and Holbrook and Krosnick (2010) flag major problems involving respondents' noncompliance with the randomized response instructions. Indeed, these authors find that randomized response produces more nonresponse and less valid estimates than a list experiment.

However, these studies do not compare randomized response estimates against the truth and, as a result, their conclusion may not be entirely warranted. Indeed, according to a comprehensive review article by Lensvelt-Mulders et al. (2005), several studies that actually validated estimates against the true sensitive information, find that the randomized response technique performs reasonably well. This conclusion is echoed by Rosenfeld, Imai, and Shapiro (2016), who find that randomized response actually outperforms other indirect methods in reducing bias and minimizing the variance of estimates.

Because we knew respondents' educational level from a prior wave of the survey, the first practice question provides a check on whether the randomized response is working. As we discuss further below, this evidence suggests that the randomized response question was implemented correctly in our survey.

4.4 Results

4.4.1 Regime support

We turn now to our results. Figure 6 compares estimates from the direct question on Putin approval to estimates from the endorsement and list experiments for all available waves of the survey. For each question type, we present both unweighted and weighted estimates. Estimates from the first wave of the survey are shown on the far left of each group in red. Note that in the second survey wave, we interviewed a fresh nationally representative cross-section (blue estimates) in addition to reinterviewing panel respondents from the first wave (green estimates). Those estimates from the second wave are shown next. Finally, estimates from the third wave are shown last, on the far right of each group, in purple. Again, both wave 1 and wave 2 were conducted in late 2021, prior

to the full-scale invasion, while wave 3 was conducted after it. A dashed horizontal line at the direct (unweighted) estimate of Putin support in wave 1 is used as a benchmark for the other estimates.

The first thing to note is that Putin's support, as measured using the direct question, stood at roughly 60 percent in late 2021, near its historic low. The endorsement experiment echoes this result, though the estimates are noisy. By contrast, the list experiment appears to indicate that Putin's true level of popularity was much lower, perhaps even below 40 percent. After the invasion, all three questioning techniques suggest that support for Putin rose.[15] Estimates of the size of this increase vary across the three methods from around 10 percentage points, according to the direct question and endorsement experiment, up to around 20 percentage points, according to the list experiment (though with wide confidence intervals).

So what is Putin's true level of support after the war? All three questioning strategies imply that Putin maintains majority support. Although imprecisely estimated, the list experiment produces the lowest of these estimates—appearing to suggest that Putin's popularity in October 2023, nearly two years into the war, was just under 60 percent. This estimate is roughly 15 percentage points lower than both the direct and endorsement estimates.

How confident are we in this estimate? First, we use the approach suggested by Aronow et al. (2015) to combining a list experiment with information from direct questioning. In effect, the combined estimator relies on the idea that respondents who self-report the sensitive attitude/trait in response to direct questioning are likely to be sincere. This is the assumption of 'no false confessions.' The combined estimator exploits this information from the direct question together with additional information from the list experiment. This results in efficiency gains (over the standard difference-in-means estimator) and less biased point estimates than under direct questioning. As can be seen on the far right of Figure 6, the Aronow et al. (2015) estimator implies that Putin's approval rating rose from around 70% in 2021 to 80% after the full-scale invasion in 2023.

[15]The the difference is imprecisely estimated and not statistically significant in the case of the endorsement experiment.

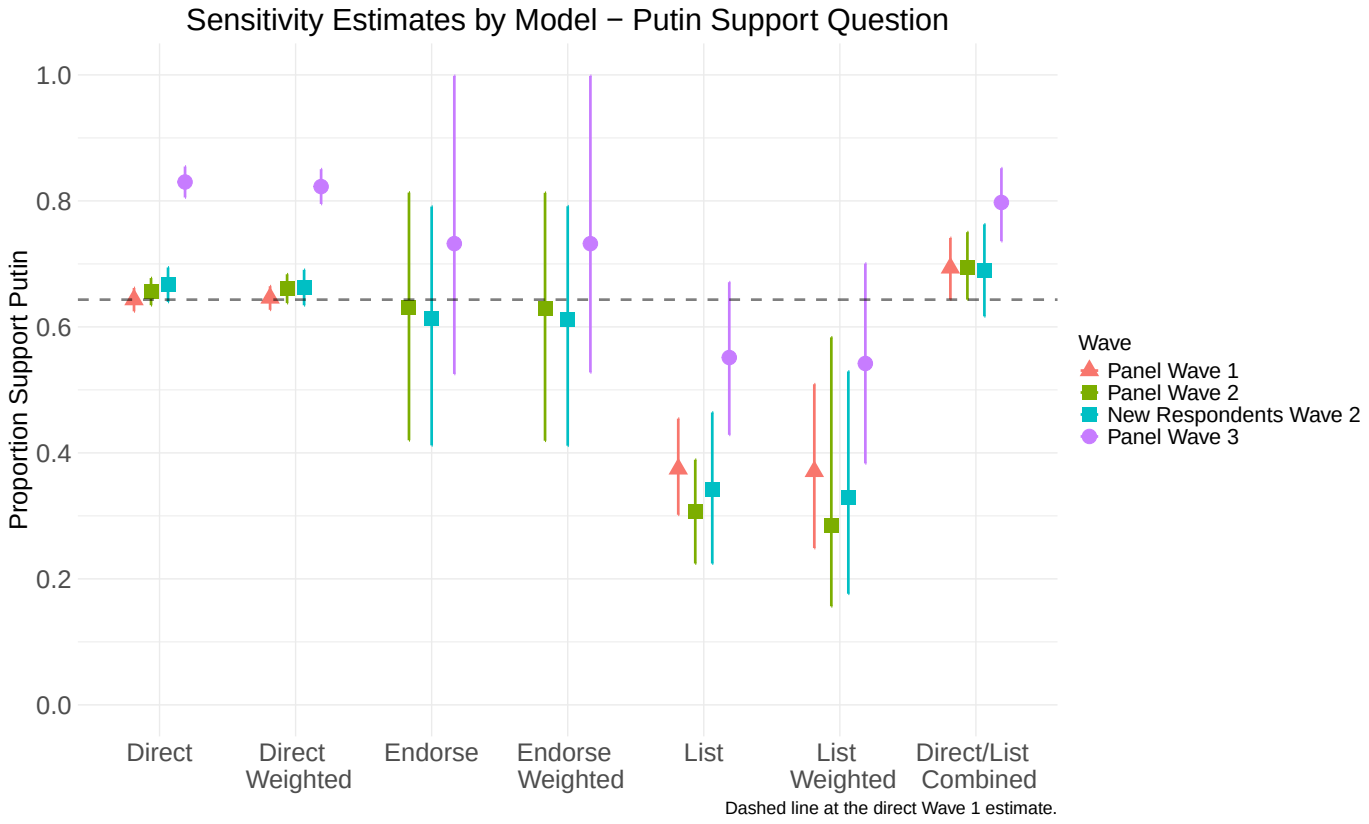

Figure 6: Estimates of Putin Support from Indirect Questioning Techniques.
This figure displays the estimated share of Russians who approve of Putin based on direct questioning, an endorsement experiment, and a list experiment (ICT) as well as the estimator suggested by Aronow et al. (2015) to combine information from direct questioning with a list experiment. Both unweighted and weighted estimates are shown. Note that there was no endorsement experiment on wave 1 of the survey.

These results are more in line with the direct and endorsement estimates and suggest little, if any, sensitivity bias.[16]

Next, we consider the problem of artificial deflation – a design effect that produces downward bias in estimates from list experiments. To address this issue, we included in the third wave of the survey a placebo list experiment. Following the advice of Frye et al. (2023), we ask both a direct question and conduct a list experiment using a non-sensitive political figure to quantify the size of the potential design effect. The rationale is that if support for a *non-sensitive* political figure appears to be lower in the list estimate than in direct questioning, this difference must be attributable to the design of the list experiment itself. For the placebo, we again ask about approval for a list of political figures. In the treatment condition, we add the name of Chinese Premier Xi Jinping, whom we argue is a non-sensitive political figure in the Russian context.[17]

Still, as Figure 7 shows, we find that the list estimate of support for Xi is 20 percentage points lower than the direct estimate. Taking this difference as our estimate of artificial deflation (i.e. deflation in the estimate of support that is not due to sensitivity but to the design of the list experiment), we conclude that Putin's true popularity after the invasion is, likely, in line with the direct and endorsement estimates: above 70 percent.

Because we did not include a placebo list in either the first or second wave of the study, and artificial inflation may vary over time (Frye et al., 2023), we are less certain about our list estimates from before the invasion. However, if we assume a similar level of artificial deflation then as now, our pre-invasion list estimates would align closely with the pre-invasion direct and endorsement estimates. These estimates place true support for Putin just above or below 60 percent.

4.4.2 War support

Last, we examine support for the war, using indirect questioning techniques to probe for evidence of dissembling. Figure 8 compares estimates from the direct question on support for continuing the war to estimates from the endorsement and list experiments.

[16] See Appendix A2 for power analyses and discussion of our ability to detect sensitivity bias under each experimental design.

[17] In particular, we do not expect that respondents would anticipate negative consequences for expressing unfavorable views of Xi.

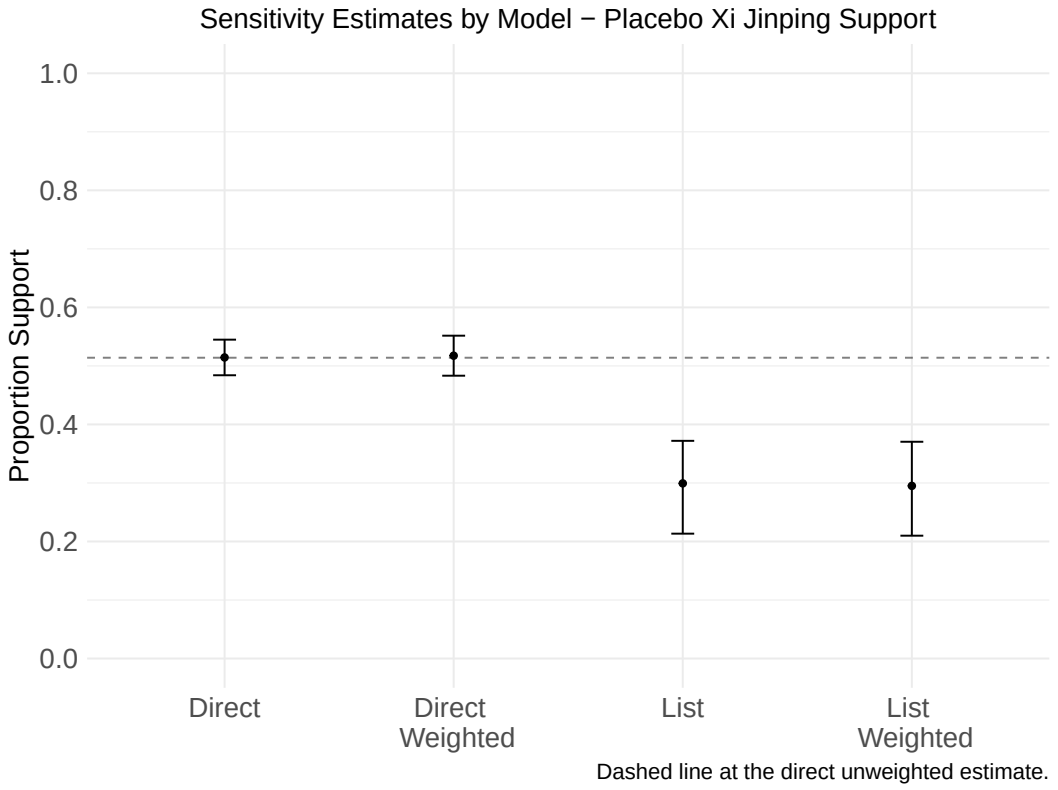

Figure 7: Estimates from Placebo List of Leaders.
This figure displays the estimated share of Russians who approve of Chinese Premier Xi Jianping, based on direct questioning and a list experiment (ICT). Both unweighted and weighted estimates are shown. The results suggest artificial deflation of around 20 percentage points in the list estimate of leader approval.

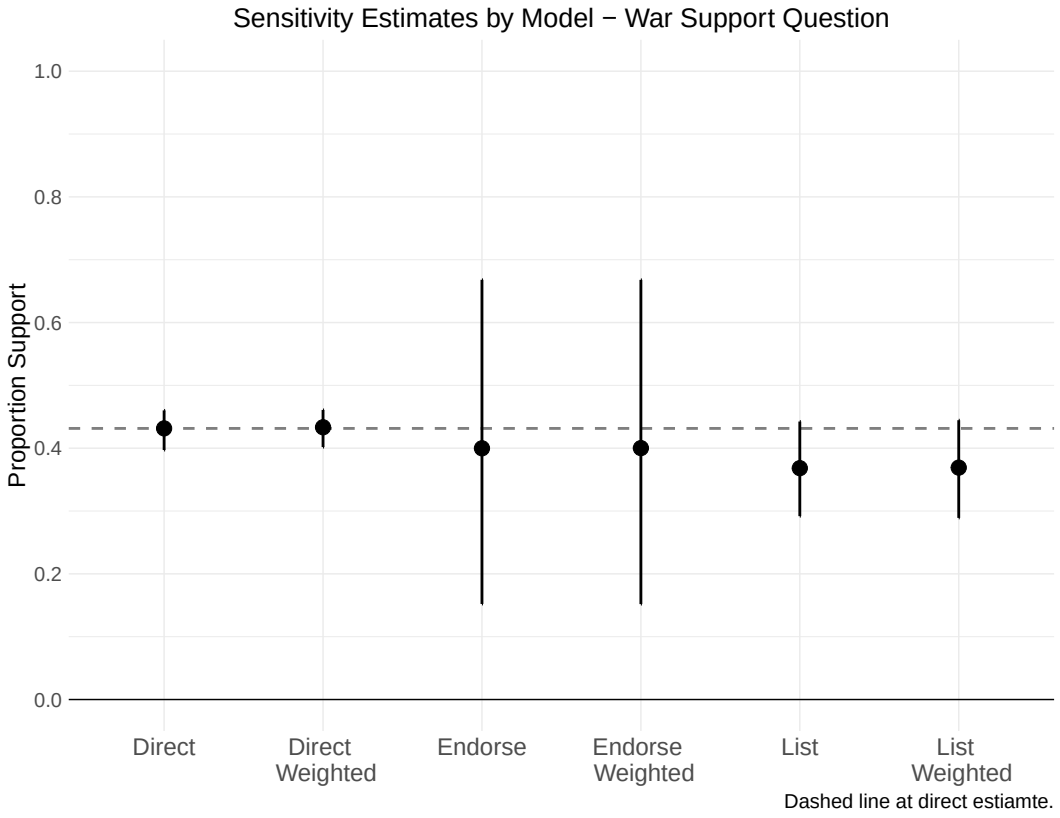

Figure 8: Estimates of War Support from Indirect Questioning Techniques.
This figure displays the estimated share of Russians who support the war in Ukraine based on direct questioning, an endorsement experiment, and a list experiment (ICT). Both unweighted and weighted estimates are shown.

Again, for each question type, we present both unweighted and weighted estimates. The dashed horizontal line indicates the direct (unweighted) estimate of war support and is used as a benchmark for the other estimates. What is immediately obvious is that estimated support for the war is very similar regardless of question type. The point estimates from the direct question are only slightly higher than the estimates from the endorsement and list experiments, and this small difference (in other words, our estimate of sensitivity bias) is statistically insignificant. Across all three questioning techniques, estimated support for the war stands at around 40 percent.

To probe the potential for artificial deflation in our list experiment estimates of war support, we again included a placebo test. This test differs from the test above insofar as the list items are statements about things the respondent may or may not do, or have done. There are also only three control items in this list, versus four above. The

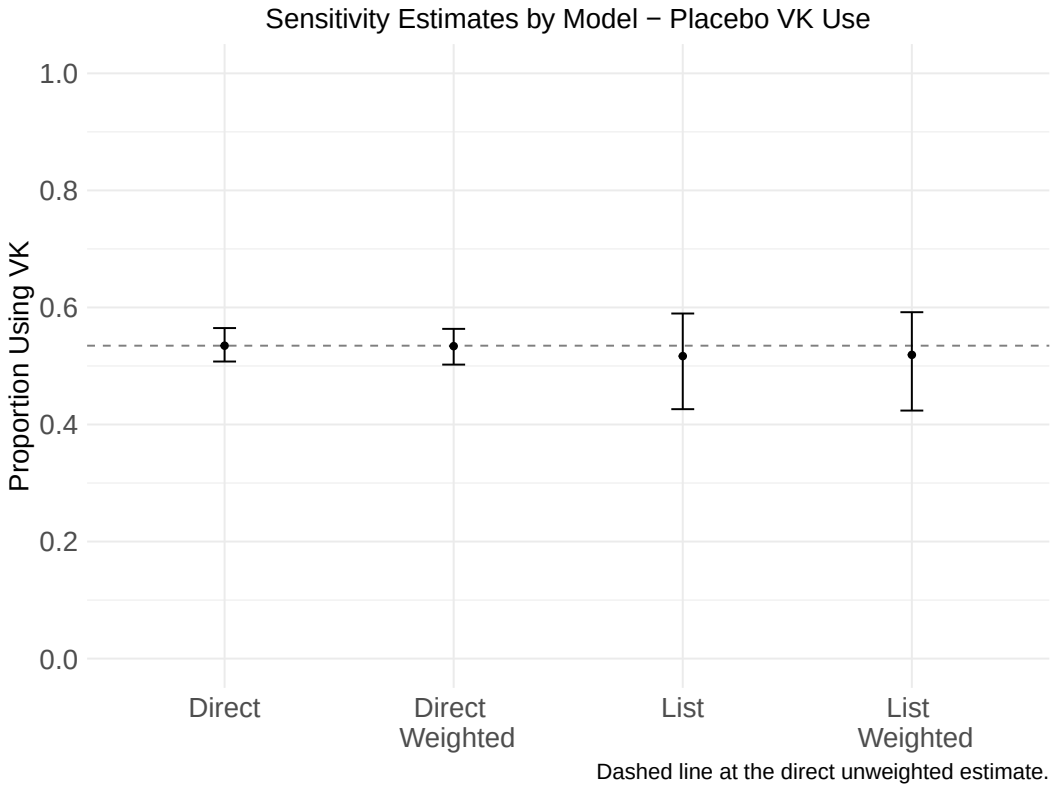

Figure 9: Estimates from Placebo List of Statements.
This figure displays the estimated share of Russians who use the social networking site VK, a nonsensitive behavior, based on direct questioning and a list experiment (ICT). Both unweighted and weighted estimates are shown. There is no evidence of artificial deflation in the list estimate.

non-sensitive placebo item is VKontakte use, which we also measure using a direct question. As shown in Figure 9, here we find no evidence of artificial deflation from the design of the list. This, thus, implies greater confidence in the accuracy of the list estimates displayed in Figure 8. In sum, these findings suggest that support for the war stands at about 40 percent and that the war, or at least continuing it, is roughly as popular as suggested by the direct question. There is little evidence, in other words, that respondents who say they support the 'special military operation' are lying.

Finally, we discuss results from the randomized response question. We present these estimates separately, since the question asked about *opposition* to the 'special military operation' in Ukraine rather than support for it, as in the list and endorsement experiments. As already discussed, if some respondents are sincerely uncertain, then the probability of support cannot be assumed to be $1 - pr(opposition)$, and vice versa. While we com-

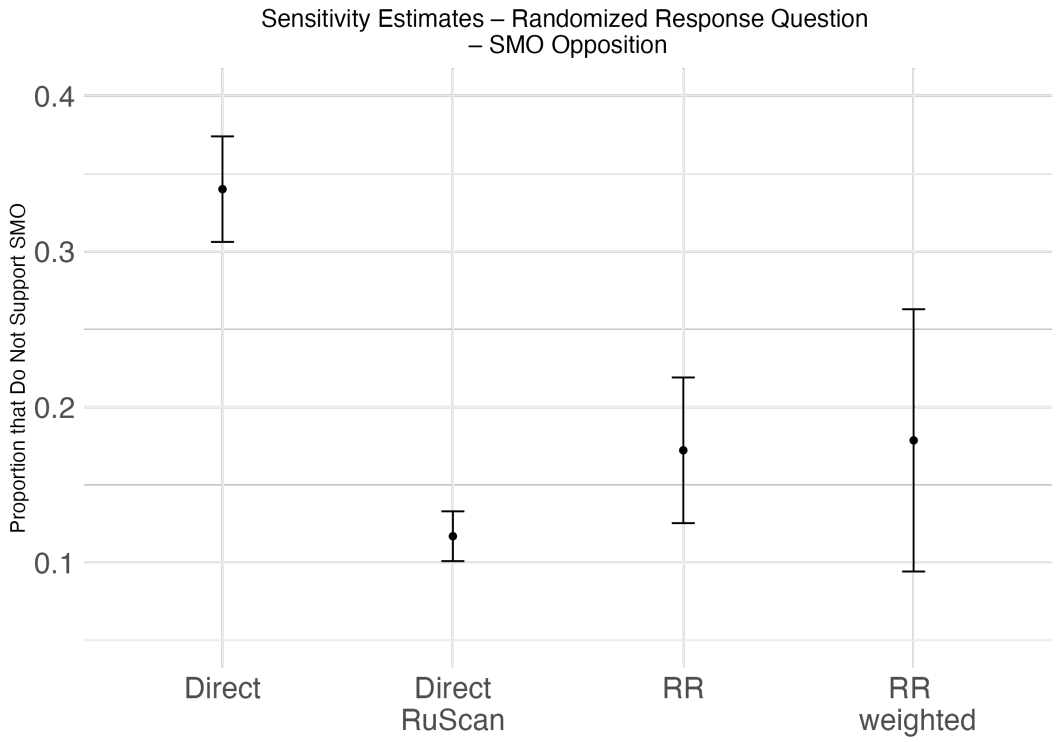

Figure 10: Estimates from Randomized Response Technique.
This figure displays the estimated share of Russians who oppose the war, using the forced-choice randomized response technique. Both unweighted and weighted estimates are shown and the results compared with direct items from two surveys.

pare our randomized response estimate of opposition to the SMO to our direct estimate, as above, we note that the randomized response question did not ask about "the continuation" of the special military operation, only the SMO itself. Our randomized response question is, thus, more similar to the question fielded around the same time by ExtremeScan, another high-quality sociological research group, that asks: "Tell me, do you support or not support Russia's special military operation in Ukraine, do you find it difficult to answer or do you not want to respond to this question?" ExtremeScan's RuScan poll was in the field September 15-26, 2023 and finds that 11.7 percent of respondents opposed the war.

Turning to Figure 10, we see that using the randomized response technique our estimate of opposition to the SMO is 17.2 percent, lower than the direct estimate from our survey, which asked about "continuation" of the war. Assuming our direct estimate is

a lower bound of actual opposition to continuing the war, these findings again imply that more Russians oppose continuing the war than describe themselves as opponents in a more general formulation of the question. At the same time, our randomized response estimate of opposition to the SMO is significantly higher than the direct estimate from ExtremeScan, though not by much, and it still suggests that opponents constitute a minority. Comparing these estimates (albeit from different polls), the implied level of preference falsification appears to be quite small, less than 6 percentage points (a finding that is significant at the $p < .05$ level). In sum, these results suggest that opposition to the war may be several percentage points higher than opinion polls suggest.

How much confidence do we have in these results? Existing studies have come to mixed conclusions about the validity of randomized response estimates. While the only existing study to validate multiple indirect techniques against ground truth (Rosenfeld, Imai, and Shapiro, 2016) finds that the randomized response technique performs very well, other scholars have raised concerns about the cognitive burden that such questions place on respondents and underscored difficulties in their implementation. To gauge directly how well respondents understand our technique and how well interviewers implement it, we included a practice round using a nonsensitive item. In this question, which also allowed respondents to familiarize themselves with the technique, we asked respondents to answer NO if they do not have higher education OR their birthday falls on an odd date, and YES otherwise. We can then compare the share of respondents who do not have higher education in our sample with the share estimated from this question using the randomized response technique. In fact, the randomized response estimate is 67 percent, while the actual share in our sample who have no higher education (i.e. had neither incomplete or complete higher or an advanced degree) was also 67 percent. This result indicates that respondents understood our question – which was asked in a way that was simple and direct – and interviewers implemented it correctly. This evidence increases our confidence in the conclusion that a small share of respondents do dissemble about their opposition to the war, but that, overall, opposition remains a minority view.

5 Conclusion

Tentatively, given the preliminary nature of all of these results, our various analyses appear to tell a fairly coherent story. They suggest that there is some preference falsification in direct questions, but that it is fairly small. It appears to have greater impact on estimates of opposition to the war than support for it. One interpretation of this evidence, alongside the evidence from our analysis of item nonresponse, which showed that risk aversion was associated with nonresponse in a binary measure of war support and with responses at the neutral midpoint of a war support thermometer, is that some opponents of the war do hide their opposition by claiming that they are uncertain about their views. However, this group is likely to be quite small, less than 10 percent overall.

References

Alyukov, Maxim. 2024. "News reception and authoritarian control in a hybrid media system: Russian TV viewers and the Russia-Ukraine conflict." *Politics* 44 (3): 400–419.

Aronow, Peter M, Alexander Coppock, Forrest W Crawford, and Donald P Green. 2015. "Combining list experiment and direct question estimates of sensitive behavior prevalence." *Journal of Survey Statistics and Methodology* 3: 43–66.

Buckley, Noah, Kyle L. Marquardt, Ora John Reuter, and Katerina Tertytchnaya. 2023. "Endogenous Popularity: How Perceptions of Support Affect the Popularity of Authoritarian Regimes." *American Political Science Review* p. 1–7.

Bullock, Will, Kosuke Imai, and Jacob N Shapiro. 2011. "Statistical analysis of endorsement experiments: Measuring support for militant groups in Pakistan." *Political Analysis* 19 (4): 363–384.

Chaiken, Shelly. 1980. "Heuristic versus Systematic Information Processing and the Use of Source versus Message Cues in Persuasion." *Journal of Personality and Social Psychology* 39 (5): 752–766.

Chapkovski, Philipp, and Max Schaub. 2022. "Solid support or secret dissent? A list experiment on preference falsification during the Russian war against Ukraine." *Research & Politics* 9 (2).

Cialdini, R.B. 1993. *Influence: The Psychology of Persuasion*. New York: Harper Collins.

Council, National Research. 2013. *Nonresponse in Social Science Surveys: A Research Agenda*. Washington, DC: The National Academies Press.

Coutts, Elisabeth, and Ben Jann. 2011. "Sensitive questions in online surveys: Experimental results for the randomized response technique (RRT) and the unmatched count technique (UCT)." *Sociological Methods & Research* 40 (1): 169–193.

Frankel, Laura Lazarus, and D. Sunshine Hillygus. 2014. "Looking Beyond Demographics: Panel Attrition in the ANES and GSS." *Political Analysis* 22 (3): 336–353.

Frye, Timothy, Scott Gehlbach, Kyle L. Marquardt, and Ora John Reuter. 2023. "Is Putin's popularity (still) real? A cautionary note on using list experiments to measure popularity in authoritarian regimes." *Post-Soviet Affairs* 39 (3): 213–222.

Hale, Henry E. 2022. "Authoritarian Rallying as Reputational Cascade? Evidence from Putin's Popularity Surge after Crimea." *American Political Science Review* 116 (2): 580–594.

Holbrook, Allyson L., and Jon A. Krosnick. 2010. "Measuring Voter Turnout by Using the Randomized Response Technique: Evidence Calling into Question the Method's Validity." *Public Opinion Quarterly* 74 (2): 328–343.

Lensvelt-Mulders, G.J.L.M., J. Hox, P.G.M. van der Heijden, and C.J.M. Maas. 2005. "Meta-analysis of randomized response research." *Sociological Methods and Research* 33: 319–348.

O'Keefe, Daniel J. 1990. *Persuasion: Theory and Research*. Sage Publications.

Petty, Richard E., and Duane T. Wegener. 1998. "Attitude Change: Multiple Roles for Persuasion Variables." In *The Handbook of Social Psychology: Volume One*, ed. Daniel T. Gilbert, Susan T. Fiske, and Gardner Lindzey. 4th ed. John Wiley and Sons.

Petty, Richard E., and John T. Cacioppo. 1986. *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*. Springer-Verlag.

Petty, Richard E., John T. Cacioppo, and David Schumann. 1983. "Central and Peripheral Routes to Advertising Effectiveness: The Moderating Role of Involvement." *Journal of Consumer Research* 10: 135–146.

Reisinger, William M., Marina Zaloznaya, and Byung-Deuk Woo. 2023. "Fear of punishment as a driver of survey misreporting and item non-response in Russia and its neighbors." *Post-Soviet Affairs* 39 (1-2): 49–59.

Rosenfeld, Bryn, Kosuke Imai, and Jacob N. Shapiro. 2016. "An Empirical Validation Study of Popular Survey Methodologies for Sensitive Questions." *American Journal of Political Science* pp. 783–802.

Shen, Xiaoxiao, and Rory Truex. 2021. "In Search of Self-Censorship." *British Journal of Political Science* 51 (4): 1672–1684.

Wood, Wendy, and Carl A. Kallgren. 1988. "Communicator Attributes and Persuasion: Recipients' Access to Attitude-Relevant Information in Memory." *Personality and Social Psychology Bulletin* 14: 172–182.

Appendix

A1 Survey Methodology

The nationally representative face-to-face panel survey on which these analyses are based is part of the long-running Russian Election Studies series. Fieldwork has been conducted by the reputable Russian polling organization the Levada Center, with support for the first two waves from a grant by the National Science Foundation. The initial surveys took place just before and after Russia's September 2021 parliamentary elections as part of a two-wave election panel. Respondents were selected using a multi-stage, stratified, probability-based design to be nationally representative of Russia's adult population. The achieved sample includes 2,501 respondents in the first wave. Data were collected August 26 - September 15, 2021. The AAPOR-1 response rate for the survey was 46%. In December, following the election, we attempted to recontact all members of the panel and interviewed an additional 1,208 respondents from a fresh nationally representative refresh sample. The fieldwork for the panel recontacts took place from December 3–22. Interviews for the fresh cross-section were conducted December 7–22. The achieved re-contact rate among panelists in the second round was 70.9%. Respondents in both waves were offered an incentive payment.

Interviews for the third (post-invasion) survey wave took place from September 14 – October 2, 2023. Only panel respondents to both the first and second waves were eligible to participate. In total, we reinterviewed 1,038 respondents, for a response rate in the third wave of 58 percent. The cooperation rate out of all respondents whom we were able to recontact in wave three was 68 percent (1,038 respondents were reinterviewed out of 1,524 who could be reached). As in the study's initial waves, respondents were offered an incentive payment of about 15 USD.

A2 Power Analyses

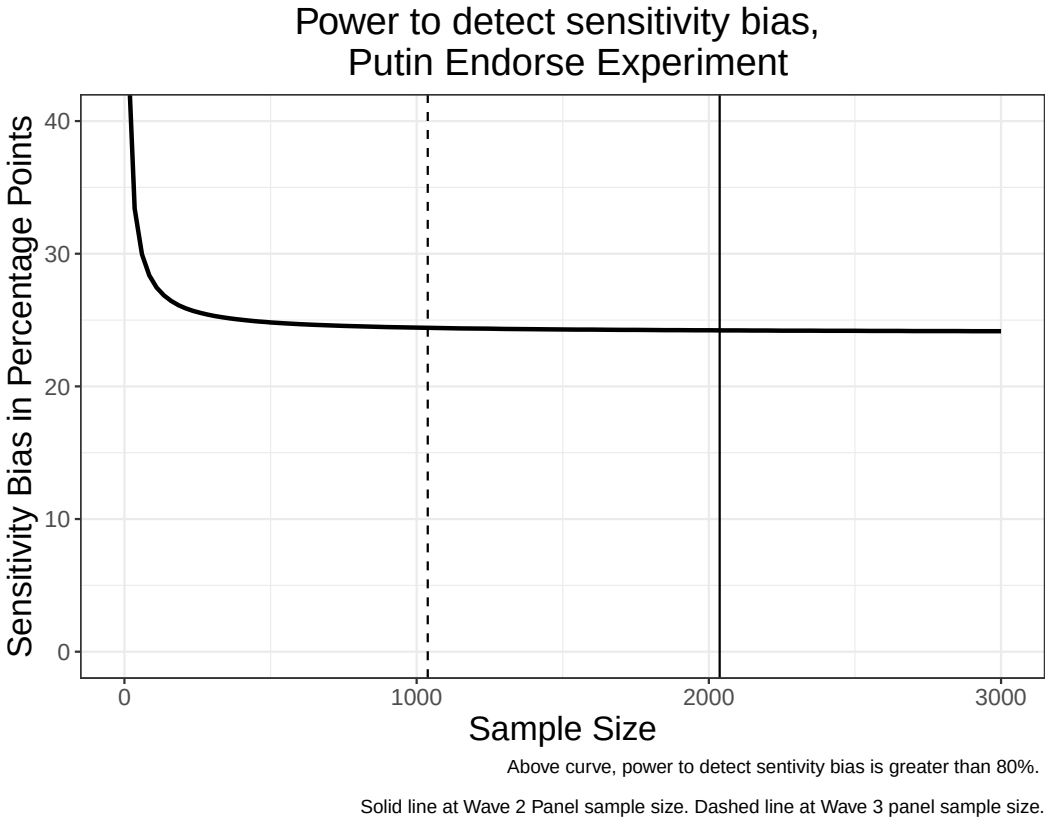

Figure A1: Power Analysis for the Putin Approval Endorsement Experiment.
This figure displays the estimated size of sensitivity bias which our experiment is powered to detect at the conventional .8 level as a function of sample size. The figure implies that our endorsement experiment is underpowered to detect sensitivity bias smaller than 24 percentage points. Thus, conservatively, with a direct estimate of approval in wave three above 80%, even if there is bias of that magnitude which we have failed to detect, Putin still likely has around 60% support.

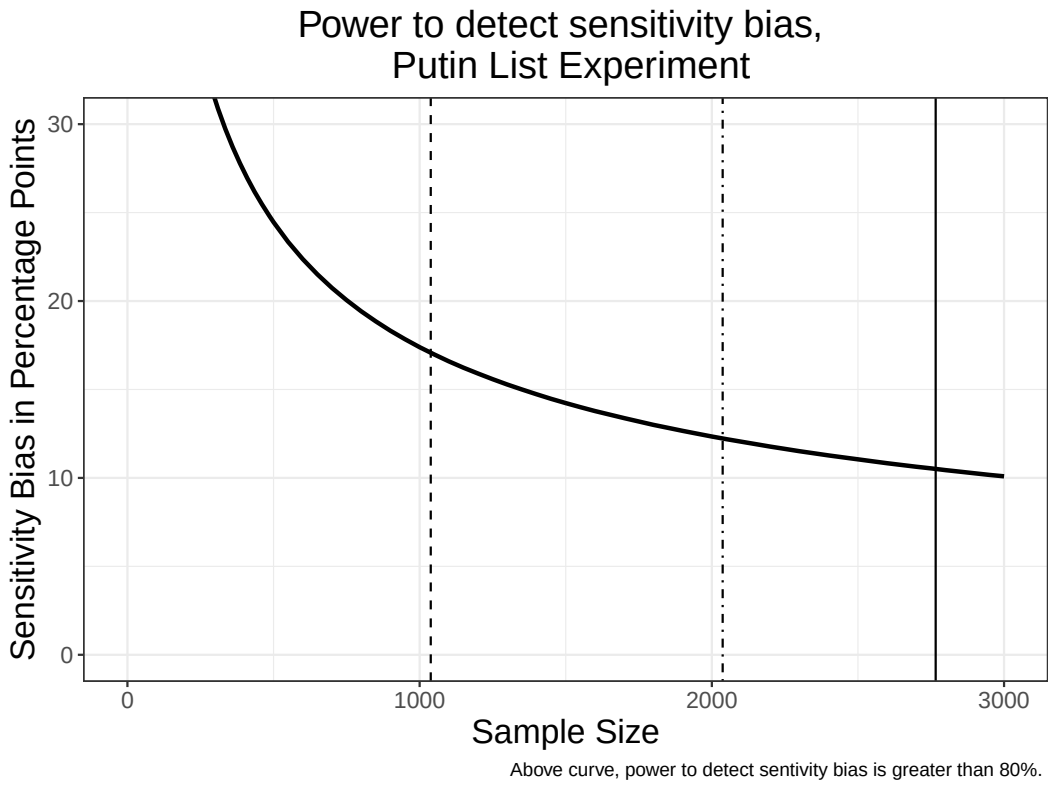

Figure A2: Power Analysis for the Putin Approval List Experiment.
This figure displays the estimated size of sensitivity bias which our experiment is powered to detect at the conventional .8 level as a function of sample size. The figure implies that our list experiment is underpowered to detect sensitivity bias smaller than 10 percentage points in wave one and 17 percentage points in wave three. Thus, conservatively, with a direct estimate of approval in wave three above 80%, even if there is bias of that magnitude which we have failed to detect, Putin still likely has more than 60% support.

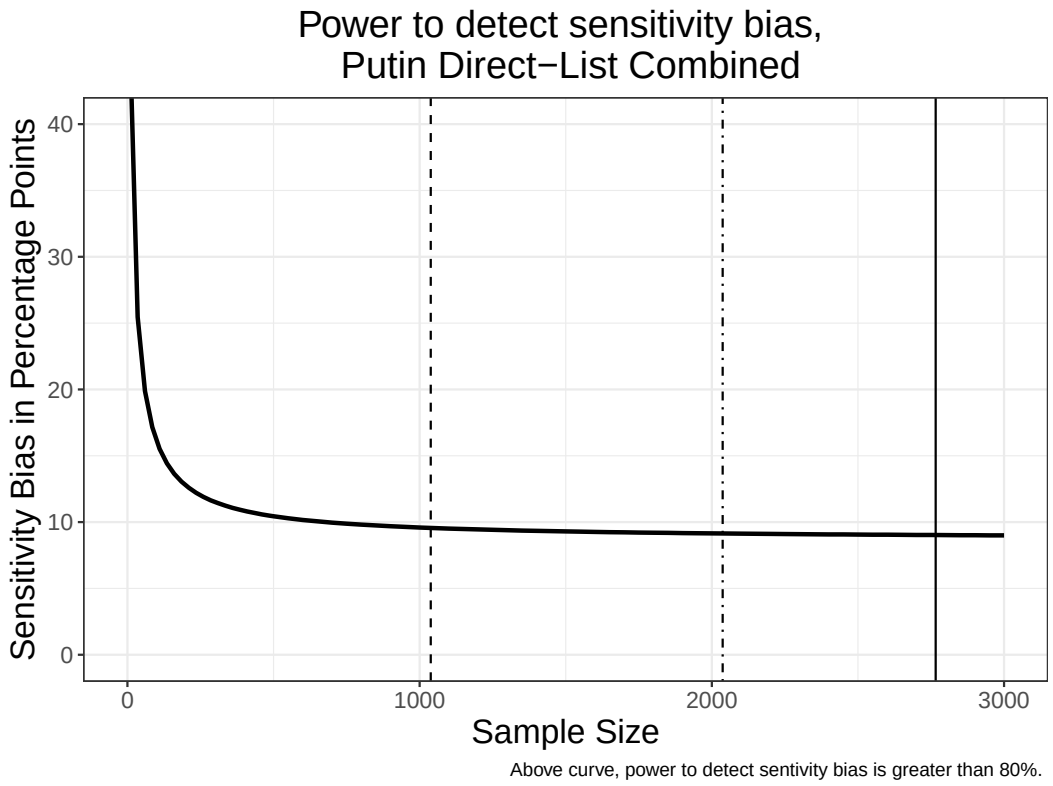

Figure A3: Power Analysis for Putin Approval, Combined Direct-List Estimator. This figure displays the estimated size of sensitivity bias which the estimator combining responses to the direct question and those of the list experiment is powered to detect at the conventional .8 level as a function of sample size. The figure implies that our combined estimator is underpowered to detect sensitivity bias smaller than 9 percentage points. Thus, conservatively, with a direct estimate of approval in wave three above 80%, even if there is bias of that magnitude which we have failed to detect, Putin still likely has more than 70% support.