The Path to a Lie*

Ciril Bosch-Rosa[†]

June 15, 2026

Abstract

Many reporting environments allow people to inspect information before submitting a report. We ask whether this search process allows people to build a path of observed evidence that makes dishonest reporting easier to justify. In an experiment embedded in the German Socio-Economic Panel Innovation Sample (SOEP-IS), participants were presented with ten boxes containing the values 0 to 9 in random order on a tablet. They could click as many boxes as they wished, but were instructed to report the value in the first box they clicked. Since participants were paid according to the number they reported, they had a monetary incentive to report a value above the first box. The first click, therefore, determined the truthful report, whereas additional clicks could only expand the set of values participants had observed before submitting their report. Exploiting the recorded click sequence, we show that participants who later misreport are especially likely to stop after a click that sets a new observed maximum. This stopping response increases with the size of the improvement. Final reports, in turn, usually remain within the values observed along the search path. These results suggest that agents can use discretionary search to create a path of observed evidence that makes a false report easier to justify.

Keywords Lying · Self-serving Justifications · Sequential Search · Survey

JEL Classification C91 · D82 · D90

*Corresponding author: Ciril Bosch-Rosa (cirilbosch@gmail.com). We would like to thank Tilman Fries for useful comments. We acknowledge financial support from the Deutsche Forschungsgemeinschaft (DFG) through the CRC TRR 190 Rationality and Competition (project number 280092119).

[†]Chair of Macroeconomics, Technische Universität Berlin & Berlin School of Economics

1 Introduction

Many decisions require people to state an outcome, summarize evidence, or justify a conclusion. These statements are often made only after the decision-maker has chosen what information to inspect and when to stop. A manager may search through performance metrics until finding one that supports a favorable review, or a researcher may examine multiple specifications before emphasizing a more favorable result. In such settings, the search path may leave the underlying truth unchanged while producing evidence that makes an untruthful report easier to justify.

This idea builds on evidence that lying is psychologically costly (Abeler et al., 2019). People are less willing to misreport when a false report feels like pure fabrication, but may be more willing to do so when the report can be tied to a favorable reference point or interpretation (Mazar et al., 2008; Shalvi et al., 2011). Prior work shows that such rationalizations can be supported by favorable counterfactuals (Shalvi et al., 2011), motivated mistakes and self-serving interpretations (Exley and Kessler, 2024; Bosch-Rosa et al., 2025a), or endogenous information acquisition (Dimant et al., 2024). Much of this evidence, however, comes from final reports, which do not reveal how they were produced. The same favorable outcome may be truthful, outright fabricated, or selected after a search process uncovers a favorable value. When the claim is false, the search path itself can become part of the justification. This shifts attention from the final response to the process that can make a false report more justifiable.

We study whether dishonest reporting is preceded by such a justificatory search using an experiment embedded in the 2020 wave of the German Socio-Economic Panel Innovation Sample (SOEP-IS) (Richter and Schupp, 2015). Participants saw ten closed boxes on a tablet screen, with values from 0 to 9 randomly assigned across boxes. They were instructed to report the value of the first box they clicked, but could inspect additional boxes before entering their report. The payment equaled the reported value in euros, regardless of whether it was truthful. The key design feature is that the first click fixed the truthful report, so additional clicks could not change what participants were instructed to report; they could only change the values participants had observed before entering their final report. Because the interface recorded the full click path, we observe the first-box truth, the values uncovered during the search, the stopping point, and the final response.

This allows us to study whether dishonest reporting is tied to favorable evidence revealed along the search path.

Our main result is that dishonest reporting leaves traces in the search path. Participants who later misreport do not simply search randomly and then enter a higher number. Instead, their search is more likely to end just after they uncover a value that improves on everything they had seen before. The larger this improvement, the stronger the stopping response. Final reports are also usually anchored in the search path: most dishonest participants report a value they actually observed, rather than jumping directly to the highest possible payoff. Taken together, these patterns suggest that dishonesty need not arise only at the moment of reporting. It can take shape along the way, as participants encounter values that make a later false report easier to justify.

Overall, this paper contributes to the literature in three ways. First, whereas much of the literature on dishonesty studies final reports, we study the process leading to dishonest reporting. The click-level data allow us to observe the search path that precedes the final response and to ask whether dishonesty is tied to that path. Second, the paper shows that endogenously acquired information can provide an observed reference point for dishonest reporting. This links work on justification in dishonesty (Mazar et al., 2008; Shalvi et al., 2011; Exley and Kessler, 2024) to work on motivated information acquisition (Dimant et al., 2024). Third, by embedding the experiment in the SOEP-IS, we study dishonest reporting and the search path that precedes it outside the convenience samples on which much of the lying literature has relied (Fosgaard, 2020).

The remainder of the paper proceeds as follows. Section 2 describes the experimental design. Section 3 presents a reduced-form model that motivates the empirical stopping tests. Section 4 reports the main empirical results. Section 5 concludes.

2 Survey Protocol and Design

The data come from the SOEP Innovation Sample (SOEP-IS) 2020 and were collected as part of the German Socio-Economic Panel (Richter and Schupp, 2015). The module was fielded as Computer-Assisted Personal Interviews (CAPI) between September and

December 2020.[1] Participants completed two incentivized reporting tasks commonly used in the experimental literature on lying. The order of the two tasks was randomized. In both tasks, payments depended solely on the participant's *report*: reporting the number $r \in \{0, \ldots, 9\}$ yielded a payment of r euros, regardless of whether the report was truthful. Participants were informed at the outset that only one task would be payoff-relevant and that a coin flip would determine which task would be paid. The full set of instructions is reproduced in Appendix C.

The main task of interest was a variant of the "box paradigm" of Gneezy et al. (2018). Participants saw 10 closed black boxes on the survey tablet, each containing an integer from 0 to 9 in random order. Clicking a box revealed its number. Participants could click as many boxes as they wished, but were instructed to report the number in the *first* box they clicked. Reports were entered privately while the interviewer was in another room. Crucially, the software recorded the full sequence of clicks, including the order in which boxes were opened and the values revealed. Participants were explicitly informed that the researcher could track all clicks, but were also reminded that all interview information was evaluated anonymously and not associated with names. The other task was a variant of the "private die-rolling paradigm" introduced by Fischbacher and Föllmi-Heusi (2013). In this case, participants rolled a fair 10-sided die under a cup and were asked to report the outcome. [2] To ensure privacy, the interviewer left the room before the roll, and participants entered their report directly on the survey tablet while alone. After both tasks were completed, the payoff was randomized, and the survey continued.

Data from this module have also been used by Bosch-Rosa et al. (2025b), who study how privacy and researchers' observability of false reports affect dishonesty and social-image costs. Their analysis focuses on comparing the private dice task and the observable box task, with particular attention to demographic heterogeneity. In contrast, we are interested in the click-level records from the box task to study how participants search before submitting a report, when they stop, and whether the final report is tied to values

[1]Despite the COVID-19 pandemic in 2020, SOEP-IS continued to do some interviews in person under strict hygiene protocols. For details, see https://www.diw.de/documents/publikationen/73/diw_01.c.818889.de/diw_ssp0986.pdf.

[2]Fischbacher and Föllmi-Heusi (2013) use a 6-sided die. We use a 10-sided die to provide a wider range of realizations (0–9) and more granular outcome distributions.

revealed along the way. These click-level records are not analyzed in Bosch-Rosa et al. (2025b). Thus, while the dice task and the privacy comparison are central to Bosch-Rosa et al. (2025b), they serve here only as complementary benchmarks for interpreting path-based dishonesty in the box task.

3 Reduced-form model

This section introduces a reduced-form model that motivates the empirical stopping tests. We do not solve the full dynamic search problem, which would require specifying participants' beliefs about unopened boxes, their costs of attention and time, and how they trade off continued search against later reporting incentives. Instead, the model isolates the payoff-side channel that is central to the experiment. Once the first box fixes the truthful report, additional search cannot change what participants are instructed to report. It can, however, reveal values that make a higher false report easier to justify. The model formalizes this idea in a simple way by allowing the cost of misreporting to depend on the best value observed along the search path.

Let $X \subset \mathbb{R}$ be the finite set of possible box values. A participant sequentially opens boxes and observes the value revealed by each opening.[3] Let $x_1 \in X$ denote the value revealed by the first box opened. The participant is instructed to report x_1, but may inspect additional boxes before submitting a response. After $k \geq 1$ first-time openings, the informational history is

$$H_k = (x_1, \dots, x_k),$$

and the best value observed so far is

$$m_k \equiv \max\{x_1, \dots, x_k\}.$$

For $k \geq 2$, define the improvement in the best value observed so far as

$$\Delta m_k \equiv \max\{0, x_k - m_{k-1}\}.$$

This variable is positive only when the current click sets a new observed maximum, and its size measures the increase in that maximum.

[3] Because repeated clicks on an already opened box reveal no new information, the model focuses on first-time box openings.

At each history H_k, the participant chooses whether to stop and submit a report or continue searching. If the participant stops, they choose a report $r \in X$. Material payoffs are normalized so that report r yields payoff r. A report above x_1 therefore increases the material payoff relative to truthful reporting, but carries a self-image cost.

The key assumption is that this self-image cost depends on what the participant has observed during the search. An upward misreport may be easier to justify when it does not exceed the best value actually seen. We capture this by allowing reports at or below m_k to receive a justification discount. Let

$$D_k \equiv D(x_1, m_k)$$

denote this discount after history H_k. Larger values of D_k make false reports less costly as long as they do not exceed the best value observed so far. Additionally, we assume that $D(x_1, m) \geq 1$ for all $m \geq x_1$, with $D(x_1, m) = 1$ corresponding to no discount, and that $D(x_1, m)$ is weakly increasing in m. The monotonicity assumption captures the idea that seeing a higher value can make a false report within the observed range feel less like pure fabrication.

We now specify the lying cost. Let $\kappa > 0$ capture baseline lying aversion. After history H_k, the cost of reporting r is

$$\psi_k(r) = \begin{cases} \dfrac{\kappa(r - x_1)_+^2}{D_k} & \text{if } r \leq m_k, \\ \\ \kappa(r - x_1)_+^2 & \text{if } r > m_k. \end{cases} \tag{1}$$

The discount applies only to reports that remain within the range of values observed during the search. Reports above m_k receive no discount, as they exceed every value observed so far. Within the discounted region, the cost remains increasing and convex in the size of the lie. Thus, bounded dishonest reports need not equal the observed maximum. Participants may instead "shade" below it, choosing some $x_1 < r < m_k$ to trade off a higher payoff against a higher lying cost.

If the participant stops after history H_k, the reporting problem depends only on the first-box value x_1 and the highest value observed so far, m_k. The payoff from stopping can therefore be written as

$$S(x_1, m_k) \equiv \max_{r \in X} r - \psi_k(r).$$

By contrast, the value of continuing can depend on more than the current maximum. Two histories with the same m_k may differ in which boxes remain unopened, and therefore in the expected value of further search. We denote the net value of continuing by $C(H_k)$. This term captures the expected benefit of opening additional boxes net of effort, time, and attention costs.[4]

The participant stops whenever

$$S(x_1, m_k) \geq C(H_k).$$

Appendix B characterizes the optimal reporting rule and shows that $S(x_1, m_k)$ is weakly increasing in m_k. Intuitively, seeing a higher value can make stopping more attractive because it expands the set of false reports that can be tied to an observed value. Solving the full dynamic stopping problem would require additional assumptions about beliefs over unopened boxes, search and attention costs, and the value of future justifying evidence. These complications are unnecessary for the empirical test. We therefore use the model to isolate the payoff-side mechanism behind the following prediction:

Hypothesis 1. Participants are more likely to stop when the current click improves the best value observed so far.

The same logic suggests that larger increases in m_k should make stopping more attractive, since they allow higher false reports to be tied to an observed value. This gives the second prediction:

Hypothesis 2. Conditional on the current click setting a new observed maximum, participants are more likely to stop when the improvement is larger.

4 Results

Of the 1,318 participants who agreed to take part in the module, 1,283 produced a valid first-box observation. The remaining 35 entered a report without clicking a box, so the

[4] Some cost of continuing is needed for stopping to be meaningful. Without such costs, additional openings would weakly dominate stopping, since they can only weakly expand the set of observed values and therefore weakly increase the observed maximum. In the data, participants typically stop well before exhausting the available search opportunities, consistent with positive costs of continuing.

first-box value is unobserved and the truthful report cannot be identified. Table 1 summarizes the main descriptive patterns. Misreporting is uncommon but non-trivial. Among participants with a valid first-box observation, 8.6% report above the first-box truth, corresponding to 110 participants.[5] Conditional on lying, misreporting is sizable, the mean lie size is 4.45 units above the truth. Search is common but far from exhaustive. Participants make 2.45 total clicks on average and open 2.16 distinct boxes. Thus, many participants search beyond the first box, but most stop well before all boxes are opened. However, search intensity varies sharply by reporting behavior. The median honest participant opens only the first box and reports it, whereas the median participant who ultimately misreports opens three distinct boxes. This suggests a link between dishonesty and the search path.[6] Since repeated openings of the same box reveal no new information, the click-level analyses below define a "click" as a first-time opening of a box.

Among the 110 participants who misreported, 14 reported above the first-box value without opening any additional boxes. The remaining 96 misreport after opening at least one additional box. Of these, 90 reported at or below the maximum observed during the search ($r \leq m_T$). Within this "bounded" group, 68 report the maximum observed ($r = m_T$), while 22 report strictly between the first-box truth and the maximum observed ($x_1 < r < m_T$). The remaining 6 report an "unbounded" value above the maximum observed ($r > m_T$). Thus, when dishonest participants search beyond the first box, most reports remain within the observed range. Only a small minority exceed the observed maximum, and a separate group misreports without any additional search.

[5] The low prevalence is consistent with broader evidence that people lie surprisingly little even when misreports are hard to detect (Abeler et al., 2019).

[6] Appendix Table A.1 reports the full distribution, separately for honest and dishonest participants.

Table 1: Descriptive statistics

	mean	sd	count
Reported value in box task [0–9]	4.949	2.869	1318
First box	4.614	2.880	1283
Misreport size (reported - first box)	0.337	1.496	1283
Misreported in box task (reported > first box)	0.086	0.280	1283
Total clicks (including revisits)	2.446	2.926	1283
Maximum value seen	5.825	2.868	1283
Dishonest subset			
Total clicks (including revisits)	4.836	4.069	110
Misreport size (reported - first box)	4.455	2.384	110
Maximum possible lie gain (9 - first box)	6.591	1.988	110

Notes: The reported-value row is based on the full module sample (N = 1,318). Rows involving the first-box value, misreporting, search, or path variables are restricted to the 1,283 participants with a valid first-box observation. The lower panel is restricted to the 110 participants who reported above the first-box value.

4.1 Favorable updates and stopping

The model predicts that participants should be more likely to stop when the current opening reveals a new best value, especially when the improvement is large. Table 2 tests this prediction at the click level by relating stopping to the improvement in the running maximum. The dependent variable $Stop_{i,k}$ equals one if participant i stops after click k and zero otherwise. The main explanatory variable is *Current improvement*, defined as

$$\Delta m_{ik} = \max\{0, x_{ik} - m_{i,k-1}\},$$

which measures how much the current click k improves the best value observed so far.

Columns (1)–(3) show that *Current improvement* positively predicts stopping. The result is robust to successively adding fixed effects for the opening position in the search path (Col. 1), the current revealed value (Col. 2), and the prior maximum and the first-box value (Col. 3). The pattern is therefore not just that participants stop after high values, but that they stop when the search path produces a new best value. Column (4) asks whether this stopping response differs by eventual reporting behavior. The positive and statistically significant interaction with *Misreported* indicates that the stopping response is stronger along paths that ultimately end in misreporting. Column (5) shows that this pattern is even larger among paths that end in reports bounded by the values observed during the search.

Table 2 focuses on timing by asking whether participants stop when the current opening improves the best value observed so far. Table 3 asks whether this result is distinct from a more general tendency to stop after opening favorable values. To do so, we use the same click-level sample and an analogous stopping specification, but replace *Current improvement* with two alternative variables. The first, *Current gap*, is defined as $\max\{0, x_{ik} - x_{i1}\}$ and measures how much the current opening exceeds the first-box value, if at all. The second, *Highest-value gap*, is defined as $\max\{0, m_{ik} - x_{i1}\}$ and measures how much the best value observed so far exceeds the first-box value.[7]

Column (1) shows that stopping is more likely when the current value is favorable relative to the first box. Column (2) shows that stopping is also more likely when the

[7] *Highest-value gap* captures how high the search path has become by opening k, whereas *Current improvement* captures only the increment at opening k. Since $m_{ik} = m_{i,k-1} + \Delta m_{ik}$ whenever the running maximum changes, the highest-value gap can be written as $(m_{i,k-1} - x_{i1}) + \Delta m_{ik}$.

Table 2: Stopping after favorable improvements in the running maximum

	(1) Stop	(2) Stop	(3) Stop	(4) Stop	(5) Stop
Current improvement	0.060***	0.045***	0.067***	0.049***	0.099***
	(0.007)	(0.009)	(0.012)	(0.013)	(0.027)
Misreported				-0.060**	
				(0.029)	
Misreported × *Current improv.*				0.052***	
				(0.014)	
Observations	1489	1489	1489	1489	314
R^2	0.140	0.155	0.175	0.184	0.341
Sample	All	All	All	All	Bounded
FE	Step	Step+value	Full	Full	Full

Notes: OLS regressions at the click level, restricted to first-time openings after the first box. The dependent variable equals one if the participant stops after that opening. *Current improvement* is defined as $\Delta m = \max\{0, x_k - m_{k-1}\}$ and measures the improvement in the running maximum generated by the current opening. "Step" denotes opening-position fixed effects, "Step+value" adds current-value fixed effects, and "Full" further adds prior-maximum and first-box fixed effects. The bounded sample in Column (5) is restricted to paths that end in a report above the first-box value and at or below the maximum value observed along the path. Standard errors are clustered at the participant level.

observed maximum is higher relative to the first-box value. Columns (3) and (4) add interactions with eventual misreporting. The interaction is positive and statistically significant for *Current gap*, but smaller and not statistically significant for *Highest-value gap*. This pattern suggests that timing matters, as search paths that end in misreporting are especially likely to stop when the current opening reveals a value above the first box. By contrast, the weaker interaction for *Highest-value gap* suggests that what matters is not simply how favorable the search history has become, but when such a favorable value is revealed.

Table 3: Alternative measures of favorability

	(1) Stop	(2) Stop	(3) Stop	(4) Stop
Current gap	0.025**		0.015	
	(0.012)		(0.013)	
Highest-value gap		0.036***		0.031***
		(0.008)		(0.009)
Misreported			-0.061*	-0.042
			(0.033)	(0.041)
Misreported × *Current gap*			0.030**	
			(0.012)	
Misreported × *Highest-value gap*				0.016
				(0.010)
Observations	1489	1489	1489	1489
R^2	0.146	0.154	0.151	0.155
FE	Yes	Yes	Yes	Yes

Notes: OLS at the click level, restricted to first-time openings after the first box. The dependent variable equals one if the participant stops after that opening. *Current gap* is $\max\{0, x\_ik - x\_i1\}$ and measures how much the current opening exceeds the first-box value. *Highest-value gap* is defined as $m_{ik} - x_{i1}$, and measures how much the highest value observed so far exceeds the first-box value. All specifications include opening position, current value, and first-box fixed effects. Standard errors clustered at the participant level.

Taken together, Tables 2 and 3 distinguish between the timing and level of favorable evidence. The first shows that stopping is especially likely when the current opening improves the best value observed so far. The second shows that stopping is also more likely when the search produces values above the first-box value, whether measured by the current opening or by the best value observed so far.

Result 1. *Participants are more likely to stop when the current opening improves the best value observed so far. This stopping response increases with the size of the improvement and is stronger along paths that ultimately end in misreporting.*

4.2 Observed evidence and final reports

The stopping evidence shows when participants end the search path. We next ask whether this timing carries over to the final report. Specifically, we test whether favorable evidence at the stopping point predicts misreporting, and whether dishonest reports remain tied to values observed along the path. Table 4 addresses the first question using linear probability models. The dependent variable, *Misreport*, equals one when the participant reports a value above the first-box value. The main explanatory variable is *Final improvement*, defined as the increase in the running maximum generated by the final opening.[8]

The results show that favorable evidence at the stopping point is associated with final misreporting. Column (1) shows a positive baseline association between *Final improvement* and misreporting. Column (2) adds first-box fixed effects and controls for search behavior.[9] The specification also controls for the number of distinct boxes opened and the highest value observed. The key comparison is therefore between participants who have the same first-box value, similar search intensity, and similar observed maxima, but who differ in the timing of favorable evidence along the search path. The results show that misreporting is more likely when this evidence arrives at the final opening. Together with

[8]For participants who opened at least one additional box after the first, *Final improvement* is defined as $\max\{0, x_{iT} - m_{i,T-1}\}$, where T denotes the final opening. For participants who only opened the first box, *Final improvement* is set to zero.

[9]First-box fixed effects compare participants with the same truthful report, which absorbs mechanical differences in the scope for upward misreporting across first-box values. The indicator *Searched* distinguishes participants who stopped after the first box from those who searched further but ended without any final improvement.

Table 4: Terminal favorable evidence and misreporting

	(1) Misreport	(2) Misreport
Final improvement	0.078***	0.045***
	(0.010)	(0.013)
Max gap		0.016*
		(0.008)
Distinct boxes		0.005
		(0.008)
Searched		0.065
		(0.043)
Observations	1283	1283
R^2	0.164	0.228
Sample	All	All
First-box FE	No	Yes

Notes: Linear probability models at the participant level. The dependent variable equals one if the participant reports above the first-box value. *Final improvement* is the improvement in the running maximum generated by the final opening before reporting; it is set to zero for participants who only opened the first box. *Max gap* is the difference between the maximum value observed along the path and the first-box value. *Distinct boxes* is the number of unique boxes opened. *Searched* equals one if the participant opened at least one additional box after the first. Column (2) includes first-box fixed effects. Robust standard errors.

Result 1, this implies that terminal improvements are associated not only with a higher likelihood of stopping, but also with a higher likelihood of misreporting.

Having linked favorable evidence at the stopping point to misreporting, we next ask whether dishonest reports remain tied to the observed search history. If the search path provides justification for dishonesty, dishonest reports should usually remain within the range of values actually observed. Among the 96 participants who misreport after searching beyond the first box, 90 report at or below the maximum value observed along the path. Figure 1 summarizes this relationship by plotting the lie size, $r - x_1$, against the gap between the observed maximum and the first-box value, $m_T - x_1$. Points on the 45-degree line correspond to reports equal to the observed maximum. The figure shows substantial bunching on this line, along with fewer shaded reports below it and very few above it. Thus, dishonest reports remain closely tied to the values observed along the path, but are not necessarily equal to the best value observed.[10]

Result 2. *Favorable evidence at the end of the search path is associated with final misreporting, and dishonest reports after additional search usually remain within the range of values actually observed.*

Taken together, Results 1 and 2 describe a path-based process of dishonesty. Participants are more likely to stop after a favorable opening. When the favorable opening occurs at the end of the search path, they are also more likely to misreport. And when they misreport after searching beyond the first box, the report usually remains within the range of values observed along the way.

[10]Among the 96 dishonest participants who search beyond the first box, 59 report a value that is both the terminal click value and the observed maximum, 9 report a previously observed maximum rather than the terminal value, 8 report the terminal value below the observed maximum, 14 choose a bounded report that is neither the terminal value nor the observed maximum, and 6 report above the observed maximum.

Figure 1: Observed evidence and shaded reports among dishonest participants

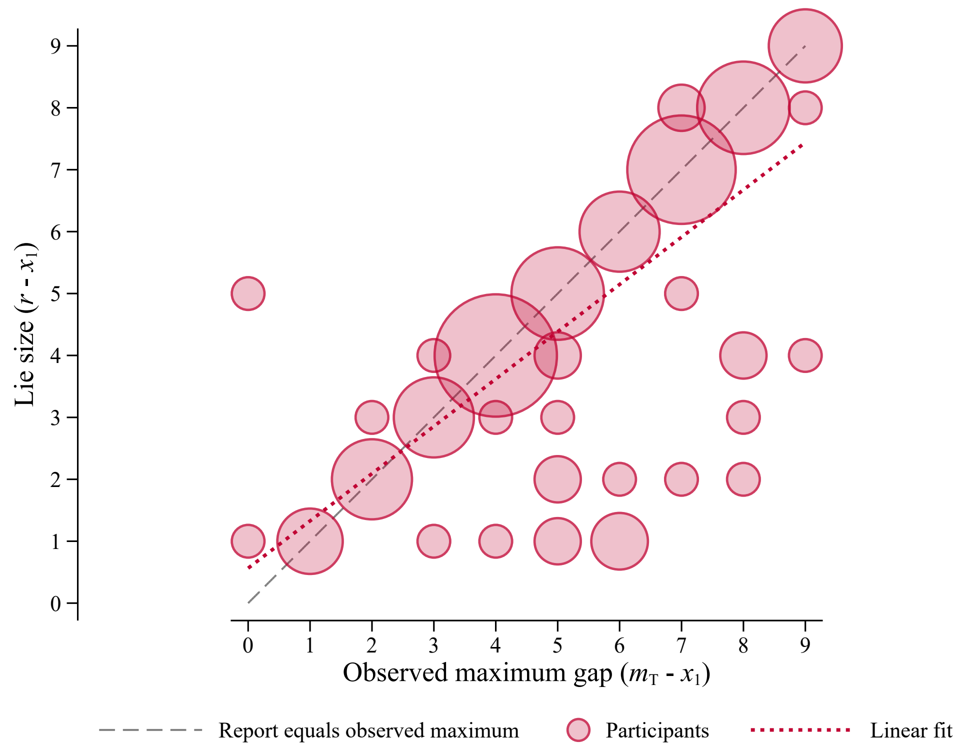

Note: The figure plots lie size, $r - x_1$, against the observed maximum gap, $m_T - x_1$, among participants who reported above the first-box value after opening at least one additional previously unopened box. Each bubble is a cell count, with larger bubbles indicating more participants. The dashed 45-degree line corresponds to reporting the observed maximum, $r = m_T$. Points on the line are exact-maximum reports, while points below the line are shaded bounded reports, $x_1 < r < m_T$.

4.3 Propensity and opportunity

In this subsection, we use the dice task as a cross-task benchmark to distinguish between two potential interpretations of the path-based patterns documented above. One interpretation is that these patterns mainly reflect a stable individual propensity to report high payoff-relevant values. A second interpretation is that the box task creates task-specific opportunities for justification, allowing participants to use values observed along the search path as reference points for the final report. The dice task is useful for this comparison because it also involves private reporting, but does not provide a search path that can generate observed reference points. Because we do not observe the true die roll, the dice task cannot identify individual lies. Instead, we use *Dice excess*, defined as the reported dice value minus 4.5, as a noisy participant-level proxy for high reporting under privacy.

Table 5 reports the main cross-task evidence. Columns (1) and (2) ask whether *Final improvement* still predicts box-task misreporting after accounting for participants' tendency to report high values under privacy (*Dice excess*) and other path features. It does. *Final improvement* remains strongly associated with box-task misreporting after controlling for *Dice excess*, while the relationship between *Dice excess* and box-task misreporting is comparatively small. Thus, path-based variation in the box task predicts misreporting beyond participants' tendency to report high values under privacy.

Columns (3) and (4) examine click-level stopping behavior. Column (3) adds *Dice excess* to the fully controlled stopping specification from Table 2. *Current improvement* remains positively associated with stopping, while the coefficient on *Dice excess* is small. This suggests that the stopping response to favorable current improvements is not merely capturing a general tendency to report high values under privacy, but reflects variation generated within the box-task search path. Column (4) then interacts *Dice excess* with *Current improvement*. The interaction is positive and statistically significant, consistent with a stronger stopping response to favorable current improvements among participants who report higher values in the dice task. This suggests that the search path creates opportunities that are used especially by participants with a greater tendency to report high payoff-relevant values under privacy.

Table 5: Cross-task robustness: dice excess as a propensity control

	Person level		Click level	
	(1)	(2)	(3)	(4)
	Misreport	Misreport	Stop	Stop
Final improvement	0.043***	0.044***		
	(0.013)	(0.013)		
Dice excess	0.005*	0.005**	0.002	-0.003
	(0.002)	(0.002)	(0.004)	(0.004)
Current improvement			0.067***	0.061***
			(0.012)	(0.012)
Current improvement × *Dice excess*				0.005**
				(0.002)
Observations	1283	1283	1489	1489
R^2	0.222	0.231	0.175	0.178
FE	No	First Box	Full	Full
Outcome mean	0.086	0.086	0.263	0.263

Notes: Columns (1)–(2) are linear probability models at the participant level. The dependent variable equals one if the participant reports above the first-box value. The path controls are *Final improvement*, *Max gap*, *Distinct boxes*, and *Searched*. *Final improvement* is the improvement in the running maximum generated by the final opening before reporting and is set to zero for participants who only opened the first box. *Max gap* is the maximum value observed along the path minus the first-box value. *Distinct boxes* is the number of unique boxes opened. *Searched* equals one if the participant opened at least one additional box after the first. Column (2) adds first-box fixed effects. Robust standard errors. Columns (3)–(4) are click-level OLS regressions, restricted to first-time openings after the first box. The dependent variable equals one if the participant stops after that opening. *Current improvement* is the improvement in the running maximum generated by the current opening. Click-level specifications absorb opening-position, current-value, prior-maximum, and first-box fixed effects. Standard errors are clustered by participant. *Dice excess* is the reported dice value minus 4.5.

Result 3. Dice excess *moderates the stopping response to favorable openings. Participants who report higher values in the dice task respond more strongly when the current opening improves the best value observed in the box task, suggesting that the search path creates opportunities that are used especially by participants with a greater tendency to report high payoff-relevant values under privacy.*

The cross-task evidence supports a propensity-and-opportunity interpretation.[11] *Dice excess* is positively related to box-task misreporting and modestly strengthens the stopping response when the current opening improves the best value observed so far. At the same time, path-based variation in the box task remains associated with both final misreporting and stopping after accounting for high reporting under privacy. The evidence, therefore, weighs against a pure propensity account and points instead to task-specific opportunities for participants to ground dishonest reports in favorable values observed along the search path.

5 Conclusion

In many reporting environments, people can inspect information before deciding what to report. We study whether this discretion allows them to build a path of observed evidence that makes dishonest reporting easier to justify. We test this idea in an incentivized module embedded in the German Socio-Economic Panel Innovation Sample (SOEP-IS). Participants saw 10 closed boxes containing the values 0 to 9 in random order and were instructed to report the value in the first box they opened. Payment depended on the reported value, and participants could open additional boxes before reporting. Additional openings could not change what participants were asked to report, but they could reveal values that made reporting a higher value easier to self-justify.

The recorded sequence of openings shows how this opportunity is used. Participants are more likely to stop immediately after an opening improves the best value observed so far, especially along paths that later end in dishonest reports within the observed range. The same timing matters for the final reporting decision. Misreporting is more

[11] Appendix Table A.2 provides supplementary descriptive evidence in the same direction. *Dice excess* is related to box-task misreporting, especially bounded misreporting, but not to whether participants search or to the number of distinct boxes they open.

likely when the last opening before the report improves the best value observed so far. When participants misreport after opening additional boxes, their reports usually remain within the values they actually observed. These patterns are not explained simply by stopping after high values or by more extensive search before misreporting. The evidence points instead to timing. Dishonest reporting is linked not only to encountering favorable evidence but also to such evidence arriving immediately before the report.

The dice task complements this interpretation by providing a proxy for the tendency to report high values under privacy. This proxy is not meaningfully related to whether participants search in the box task or to the number of distinct boxes they open. It matters instead when the search path produces favorable observed evidence. The pattern suggests that dishonest reporting is associated not only with a general tendency to report high values, but also with opportunities created by the search path. The evidence does not separately identify whether participants search in order to find justifying evidence or instead follow a broader satisficing rule, stopping once they encounter a value they are willing to report.

Overall, our results suggest that dishonesty can take shape before a final report is entered. Once the truthful report is fixed, additional search cannot change what participants are asked to report. It can, however, change what they have observed. This matters because a false report may then be tied to evidence revealed along the way. Our setting is stylized, but it isolates a broader mechanism. Discretionary search can matter not only by changing what people know, but also by changing what they can point to when reporting.

References

Abeler, J., D. Nosenzo, and C. Raymond (2019): "Preferences for Truth-Telling," *Econometrica*, 87, 1115–1153. Cited on pages 2 and 8.

Bosch-Rosa, C., D. Gietl, and F. Heinemann (2025a): "Risk taking under limited liability and moral hazard: Quantifying the role of motivated beliefs," *Management Science*, 71, 976–991. Cited on page 2.

Bosch-Rosa, C., D. Nosenzo, and L. Neyse (2025b): "Observability, Honesty, and the Social Image Costs of Lying," *Available at SSRN 5080148*. Cited on pages 4 and 5.

Dimant, E., F. Galeotti, and M. C. Villeval (2024): "Motivated information acquisition and social norm formation," *European Economic Review*, 167, 104778. Cited on pages 2 and 3.

Exley, C. L. and J. B. Kessler (2024): "Motivated Errors," *American Economic Review*, 114, 961–987. Cited on pages 2 and 3.

Fischbacher, U. and F. Föllmi-Heusi (2013): "Lies in Disguise—An Experimental Study on Cheating," *Journal of the European Economic Association*, 11, 525–547. Cited on page 4.

Fosgaard, T. R. (2020): "Students Cheat More: Comparing the Dishonesty of a Student Sample and a Representative Sample in the Laboratory," *The Scandinavian Journal of Economics*, 122, 257–279. Cited on page 3.

Gneezy, U., A. Kajackaite, and J. Sobel (2018): "Lying Aversion and the Size of the Lie," *American Economic Review*, 108, 419–453. Cited on page 4.

Mazar, N., O. Amir, and D. Ariely (2008): "The Dishonesty of Honest People: A Theory of Self-Concept Maintenance," *Journal of Marketing Research*, 45, 633–644. Cited on pages 2 and 3.

Richter, D. and J. Schupp (2015): "The SOEP Innovation Sample (SOEP IS)," *Journal of Contextual Economics – Schmollers Jahrbuch*, 135, 389–399. Cited on pages 2 and 3.

Shalvi, S., J. Dana, M. J. J. Handgraaf, and C. K. W. De Dreu (2011): "Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior," *Organizational Behavior and Human Decision Processes*, 115, 181–190. Cited on pages 2 and 3.

A Additional tables

Table A.1: Distribution of search intensity, by misreporting status

	Honest		Dishonest	
	Distinct	Clicks	Distinct	Clicks
25th percentile	1	1	2	2
Median	1	1	3	3
75th percentile	2	2	6	7
90th percentile	5	5	8.5	10.5
95th percentile	8	9	10	15
99th percentile	10	15	10	15
Maximum	10	15	10	15
Mean	1.99	2.22	4	4.84
Standard deviation	2.15	2.69	2.73	4.07

Notes: Person-level search intensity in the valid first-box sample (Honest N = 1,173; Dishonest N = 110). *Distinct* is the number of unique boxes opened on the search path. *Clicks* counts every box opening, including revisits.

Table A.2: Dice-task reports and box-task behavior

	(1) Dice excess	(2) Dice excess	(3) Searched	(4) Distinct boxes	(5) Final improvement
Box misreport	0.645**				
	(0.266)				
Bounded misreport		0.778***			
		(0.283)			
Other misreport		0.047			
		(0.661)			
Dice excess			0.004	-0.000	0.037***
			(0.005)	(0.023)	(0.014)
Observations	1283	1283	1283	1283	1283
R^2	0.004	0.005	0.049	0.025	0.100
Task order	Yes	Yes	Yes	Yes	Yes
First-box FE	No	No	Yes	Yes	Yes
Outcome mean	0.759	0.759	0.306	2.161	0.483

Notes: Linear regressions at the participant level. Dice excess is the reported dice value minus 4.5, the expected value under truthful reporting. *Box misreport* equals one if the participant reports above the first-box value in the box task. *Bounded misreport* equals one if the participant reports above the first-box value and at or below the maximum value observed along the path. *Other misreport* collects box-task misreports that exceed the observed maximum or occur without additional search. *Searched* equals one if the participant opened at least one additional box after the first. *Distinct boxes* is the number of unique boxes opened. *Final improvement* is the improvement in the running maximum generated by the final opening before reporting and is set to zero for participants who only opened the first box. Cols. (3)–(5) include first-box fixed effects; Cols. (1)–(2) do not, since the first-box value is uninformative for dice reports. Robust standard errors.

B Reporting characterization and monotonicity of the stopping payoff

This appendix formalizes the reduced-form model in Section 3. The reporting problem is static conditional on the realized state, so the characterization depends only on (x_1, m_k).

B.1 Optimal reporting rule

Fix a realized state after k openings and write

$$D_k \equiv D(x_1, m_k).$$

Using the cost function in (1), the conditional reporting problem is

$$\max_{r \in X} \widetilde{U}_k(r), \qquad \widetilde{U}_k(r) \equiv r - \psi_k(r).$$

Lemma B.1. *Every optimal report satisfies $r \geq x_1$.*

Proof. If $r < x_1$, then $(r - x_1)_+ = 0$ and the lying cost is zero, so $\widetilde{U}_k(r) = r$. The same holds at $r = x_1$, giving $\widetilde{U}_k(x_1) = x_1$. Hence

$$\widetilde{U}_k(r) < \widetilde{U}_k(x_1),$$

and no report below x_1 can be optimal. □

By Lemma B.1, it is enough to consider reports weakly above x_1. Define the bounded and unbounded regions as

$$X_k^B \equiv \{r \in X : x_1 \leq r \leq m_k\}, \qquad X_k^U \equiv \{r \in X : r > m_k\}.$$

Proposition B.1. *Fix a state (x_1, m_k).*

On the bounded region X_k^B, the objective is

$$\widetilde{U}_k^B(r) = r - \kappa \frac{(r - x_1)^2}{D_k}.$$

The unconstrained continuous maximizer of this quadratic is

$$r_k^{B,c} = x_1 + \frac{D_k}{2\kappa}.$$

Imposing the upper bound $r \le m_k$ gives

$$\widehat{r}_k^B = \min\left\{m_k,\ x_1 + \frac{D_k}{2\kappa}\right\}.$$

The bounded-region maximizers are the feasible support point or points in X_k^B closest to $\widehat{r}_k^B$.

On the unbounded region X_k^U, the objective is

$$\widetilde{U}_k^U(r) = r - \kappa(r - x_1)^2.$$

The unconstrained continuous maximizer of this quadratic is

$$r_k^{U,c} = x_1 + \frac{1}{2\kappa}.$$

If $X_k^U \neq \emptyset$, the unbounded-region maximizers are the feasible support point or points in X_k^U closest to $r_k^{U,c}$. If $X_k^U = \emptyset$, no unbounded report is feasible.

Let

$$V_k^B \equiv \max_{r \in X_k^B} \widetilde{U}_k^B(r), \qquad V_k^U \equiv \begin{cases} \max_{r \in X_k^U} \widetilde{U}_k^U(r) & \textit{if } X_k^U \neq \emptyset, \\ -\infty & \textit{if } X_k^U = \emptyset. \end{cases}$$

The global optimum is obtained by comparing V_k^B and V_k^U.

Proof. On the bounded region,

$$\widetilde{U}_k^B(r) = r - \kappa\frac{(r - x_1)^2}{D_k} = x_1 + \frac{D_k}{4\kappa} - \frac{\kappa}{D_k}\left(r - x_1 - \frac{D_k}{2\kappa}\right)^2.$$

Thus the continuous maximizer is $x_1 + D_k/(2\kappa)$. Since $D_k \ge 1$ and $\kappa > 0$, this value is weakly above x_1. Imposing $r \le m_k$ gives

$$\widehat{r}_k^B = \min\left\{m_k,\ x_1 + \frac{D_k}{2\kappa}\right\}.$$

On the finite set X_k^B, the maximizers are the feasible support point or points closest to this value.

On the unbounded region,

$$\widetilde{U}_k^U(r) = r - \kappa(r - x_1)^2 = x_1 + \frac{1}{4\kappa} - \kappa\left(r - x_1 - \frac{1}{2\kappa}\right)^2.$$

Thus the unconstrained continuous maximizer is $x_1 + 1/(2\kappa)$. If X_k^U is nonempty, the maximizers on the finite set X_k^U are the feasible support point or points closest to this value. If X_k^U is empty, no unbounded report is feasible.

By Lemma B.1, no report below x_1 can be optimal. Every remaining feasible report lies either in X_k^B or in X_k^U. The global optimum is therefore obtained by comparing the best payoff attainable in the bounded region, V_k^B, with the best payoff attainable in the unbounded region, V_k^U. If X_k^U is empty, then $V_k^U = -\infty$ and the optimum is bounded. □

For fixed x_1, m_k, and κ, the bounded continuous optimum $\widehat{r}_k^B$ is weakly increasing in D_k. By contrast, the payoff from reports above the running maximum is unaffected by the justification factor. For fixed x_1, m_k, and κ, the bounded continuous optimum $\widehat{r}_k^B$ is weakly increasing in D_k. By contrast, the payoff from reports above the running maximum is unaffected by the justification factor. When

$$x_1 + \frac{D_k}{2\kappa} < m_k,$$

the bounded continuous optimum lies below the running maximum. In the finite-support problem, the optimal bounded report is then the feasible support point closest to this interior target, so it may involve shading below m_k. Otherwise, $\widehat{r}_k^B = m_k$, and the bounded optimum bunches at the running maximum.

In words, the reporting problem reduces to a comparison between two candidates. One is the best feasible report that remains within the observed search path. The other is the best feasible report that exceeds it. The first benefits from the justification discount, whereas the second does not.

B.2 Monotonicity of the stopping payoff

For the monotonicity result, write the history-specific objects from the main text as functions of a generic running maximum $m \geq x_1$. Define

$$S(x_1, m) \equiv \max_{r \in X} \{r - \psi(r; x_1, m)\}.$$

Assume that $D(x_1, m) \geq 1$ for all $m \geq x_1$, and that $D(x_1, m)$ is weakly increasing in m.

Proposition B.2. *Fix x_1 and two states $m, m' \in X$ with*

$$x_1 \leq m \leq m'.$$

Then

$$S(x_1, m') \geq S(x_1, m).$$

Proof. Fix any report $r \in X$. We show that the payoff from this fixed report is weakly higher at m' than at m.

First suppose that $r \le m$. Then $r \le m'$ as well, so the discounted branch of the cost function applies in both states. Since $D(x_1, m') \ge D(x_1, m)$ and $(r - x_1)_+^2 \ge 0$, the lying cost is weakly lower at m' than at m.

Second suppose that $m < r \le m'$. Then the report is above the running maximum at m, but not above the running maximum at m'. Its cost at m is therefore

$$\kappa(r - x_1)_+^2,$$

whereas its cost at m' is

$$\frac{\kappa(r - x_1)_+^2}{D(x_1, m')}.$$

Because $D(x_1, m') \ge 1$, the cost is weakly lower at m'.

Finally suppose that $r > m'$. Then the report is above the running maximum in both states, so the unbounded branch applies in both states and the lying cost is unchanged.

Thus, for every fixed report $r \in X$,

$$r - \psi(r; x_1, m') \ge r - \psi(r; x_1, m).$$

Taking maxima over r gives

$$S(x_1, m') = \max_{r \in X}\{r - \psi(r; x_1, m')\} \ge \max_{r \in X}\{r - \psi(r; x_1, m)\} = S(x_1, m).$$

□

Proposition B.2 is the only general stopping-payoff result established by the model. It shows that a higher running maximum weakly raises the payoff from stopping. In particular, once the running maximum rises above x_1, reports in the interval $x_1 < r \le m$ become eligible for the justification discount. The bounded component of the stopping payoff may then increase strictly if a history-consistent dishonest report yields a payoff above that of truthful reporting. Without further structure on the continuation value $C(H_k)$, however, the model does not yield a global comparative static for stopping behavior. This is why the empirical analysis treats favorable updates in the running maximum as payoff-motivated reduced-form predictions rather than as direct implications of a fully solved stopping problem.

C Instructions

Below we transcribe the instructions that participants saw on their screens. We also include interviewers' instructions in *italics*. Importantly, participants did not see the interviewer's instructions.

All of these instructions come from the module documentation which can be found online at the SOEP-IS webpage.

Intro

Now I would like to invite you to participate in a short and simple game. The task consists of two parts. What the two parts are about, you will learn immediately before. In both parts, you have a chance to win real money. But only one of the two parts will be selected for the final payout. You will flip a coin for it. If the coin shows "tails" after the toss, you will get the result for part 1. If the coin shows "heads", you will get the result for part 2. All information from the game, like all other information from this interview, will be evaluated anonymously only and will not be associated with your name. This is the end of the introduction.

Interviewer: Please do not inform the target person in advance about the content of the two games! As the interviewer, make sure that the target person does not see the dice cup and dice before the dice game begins. Please do not give the dice cup and dice to the target person until the target person asks you for it!

Q490

Would you like to play the game?

Interviewer: Before starting the module: As the interviewer, please try to make sure that no other interviewee is within earshot. As the interviewer, also try to leave the room or turn away so that the target person can answer the questions alone. It is important that the target person can complete the following tasks in complete privacy. If the target person agrees to play the game, please hand over the laptop to the target person for self-completion. After the experiment, the target person will be asked to hand over the laptop to you and continue the interview as regular.

- Yes

- No

Q491

Interviewer: Please hand over the laptop to the interviewee for self-completion.

[assume the randomization has participants start with the unobservable task]

Q493

The interviewer has just given you a 10-sided die and a dice cup. Now please roll the dice as many times as you like using the dice cup. Caution: Make sure that the interviewer does not see the result of your roll. Please remember the result of your first roll and enter it on the screen. How much you can earn in this part, if this part of the task is selected for payout, depends on the number you write down here. A higher number means a higher payout. You can see the exact payout for each reported number in the table below. The table will remain on the screen until you enter your number.

Note: Please have your interviewer hand you the dice cup and dice. Please make sure that no one can observe you playing the game.

[Participants saw a table reflecting the payoffs for all values they could report]

Q494

In this part, you will see 10 black boxes on your screen. There are numbers behind the black boxes. The numbers behind the black boxes are 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9. The numbers are in random order, so the first box does not necessarily contain 1, the second does not necessarily contain 2, and so on. As soon as you click on a black box, the number hidden behind it is displayed. You can click on as many black boxes as you like. The laptop records how many and which boxes you click. Please remember the number behind the first black box you clicked and enter it on the screen. How much you earn in this part, if this part of the task is selected for payout, depends on the number you enter. A higher number means a higher payout. You can see the exact payout for each number reported in the table below. The table will remain on the screen until you enter your number.

[Participants saw a table reflecting the payoffs for all values they could report]

Q495

You have reached the end of the task. Please hand the laptop back to your interviewer(s) and continue the survey as usual.

Q496_0

Now we will determine whether the first or second part of this task will be paid out. To do this, I ask you to flip a coin in a clearly visible position so that I can note down the result. Ask the target person to toss a coin in plain sight. Note whether "tails" or "heads" is on top.

Interviewer: Ask the target person to toss a coin in plain sight. Note whether "tails" or "heads" is on top.

Q496

As you can see, [PROG: Show result from IZAHLAUSZ01] is on top. Therefore, you will now receive the payoff from [PROG: please show: Part 1 (if IZAHLAUSZ01=1); Part 2 (if IZAHLAUSZ01=2)]. You have indicated that in [PROG: please show: Part 1 (if IZAHLAUSZ01=1); Part 2 (if IZAHLAUSZ01=2)] the first number was [PROG: please enter the number the participant entered based on their split group (SPLITZAHL) in either Part 1 or Part 2]. Therefore, you will now receive euros from me [PROG: reported number = euro amount].

Interviewer: Please pay the target person the appropriate amount.