Online Appendix

Overinference from Weak Signals and Underinference from Strong Signals

Ned Augenblick, Eben Lazarus, and Michael Thaler

Contents

A. Proofs and Additional Theoretical Discussion

A.1. Proofs for Section II.A

Proof of Proposition 1. Fix an arbitrary direction s_d. Given Assumption 3, we can write $\hat{\mathbb{S}}(\hat{s}) = \alpha(e)e + (1-\alpha(e))\hat{\mathbb{S}}(s_d)$ for some $\alpha(e) \in (0,1)$, where $\alpha(\cdot)$ may depend on s_d. We want to characterize

$$\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) = \mathbb{E}\Big[(1-\alpha(e))\hat{\mathbb{S}}(s_d) + \alpha(e)e \,\Big|\, s\Big] - \mathbb{S}(s). \tag{A-1}$$

For notational convenience, assume a continuous space of estimates e (in which $p(e|s)$ is a probability density function with support $E \subseteq \mathbb{R}$).[1] From Assumption 1, $\mathbb{E}[e|s] = \mathbb{S}(s)$, so $\mathbb{S}(s) = \int_E e\, p(e|s)\, de$, with $p(e|s)$ non-degenerate. Using this in (A-1),

$$\begin{aligned}\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) &= \left[\int_E \Big((1-\alpha(e))\hat{\mathbb{S}}(s_d) + \alpha(e)e\Big)\, p(e|s)\, de\right] - \mathbb{S}(s) \\ &= \int_E (1-\alpha(e))(\hat{\mathbb{S}}(s_d) - e)\, p(e|s)\, de. \end{aligned} \tag{A-2}$$

Denote $g(e) \equiv (1-\alpha(e))(\hat{\mathbb{S}}(s_d) - e)$. For the first term in $g(e)$, Assumption 3 gives that $1-\alpha(e) > 0$. For the second term, $\hat{\mathbb{S}}(s_d) - e$ crosses 0 exactly once for $e \in \mathbb{R}$: it is positive for $e < \hat{\mathbb{S}}(s_d)$, and negative for $e > \hat{\mathbb{S}}(s_d)$. We thus have that

$$\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) = \int g(e)\, p(e|s)\, de = \mathbb{E}[g(e)|s] = \mathbb{E}[g(e)|s_d, \mathbb{S}], \tag{A-3}$$

where $g(e)$ is strictly single-crossing from above and where $p(e|s) = p(e|s_d, \mathbb{S})$ has the strict MLRP in $\mathbb{S}$, from Assumption 1(b). By the variation diminishing property of Karlin (1968),[2] the expectation of a strictly single-crossing function with respect to an MLRP distribution is also strictly single-crossing, with the same arrangement of signs as the function (here, positive and then negative). That is, if $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) = 0$ at $\mathbb{S}(s) = \mathbb{S}^*$, then there is overreaction ($\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) > 0$) for $\mathbb{S}(s) < \mathbb{S}^*$ and underreaction ($\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) < 0$) for $\mathbb{S}(s) > \mathbb{S}^*$.

Further, this switching point $\mathbb{S}^*$ must exist and lie within the range of feasible values

[1] The steps in the proof carry through for discrete e when replacing integrals with sums and adjusting straightforwardly (though tediously) for discontinuities.

[2] See Karlin's Theorem 3.1 of Chapter 5, or Gollier (2001, Proposition 16) for a textbook reference based on the generalization of Athey (2002, Theorem 2). These results are typically stated for a function that is single-crossing from below (SCB); in our case, one can define the SCB function $\tilde{g}(e) \equiv -g(e)$ and then take $\mathbb{S}(s) - \mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] = \int \tilde{g}(e)\, p(e|s)\, de$, and all the statements carry through with signs changed appropriately. Note also that (A-3) can be restated, suppressing dependence on the arbitrary and fixed s_d, as $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|\mathbb{S}] - \mathbb{S} = \mathbb{E}[g(e)|\mathbb{S}]$, and it is this expression to which we apply Karlin's result. (Note that all references in this Online Appendix are listed in the reference list in the main text.)

$\mathbb{S}(s) \in [\min_{s_m}\mathbb{S}(s_d, s_m), \max_{s_m}\mathbb{S}(s_d, s_m)]$. To see this, consider the case $\mathbb{S}(s) = \min_{s_m}\mathbb{S}(s_d, s_m)$. By Assumption 2, $\hat{\mathbb{S}}(s_d) > \mathbb{S}(s)$ in this case, while $\mathbb{E}[e|\mathbb{S}] = \mathbb{S}(s)$ by Assumption 1. Thus $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] = \mathbb{E}[\alpha(e)e+(1-\alpha(e))\hat{\mathbb{S}}(s_d)|s]$ must satisfy $\mathbb{S}(s) < \mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] < \hat{\mathbb{S}}(s_d)$ by Assumption 3, where the lower bound obtains from $\alpha(e) \to 1$ and $\mathbb{E}[e|\mathbb{S}] = \mathbb{S}(s)$ (and the upper bound obtains from $\alpha(e) \to 0$). Thus $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) > 0$ at this minimal $\mathbb{S}(s)$. The same argument gives that $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) < 0$ at the maximal $\mathbb{S}(s)$. The intermediate value theorem then gives that there is such a switching point $\mathbb{S}^* \in (\min_{s_m}\mathbb{S}(s_d, s_m), \max_{s_m}\mathbb{S}(s_d, s_m))$ at which $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s) = \mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}^* = 0$, and the single-crossing result above guarantees its uniqueness, completing the proof. $\square$

Derivation of Monotonicity Results. As at the end of Section II.A, under Assumptions 1–3, it is not necessarily the case that a person's expected signal strength $\hat{\mathbb{S}}(\hat{s})$ is monotonic in e or that the amount of over- or underreaction $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s)$ is monotonic in $\mathbb{S}(s)$. For conditions under which these additional monotonicity results hold, we again use Assumption 3 to write $\hat{\mathbb{S}}(\hat{s}) = \alpha(e)e + (1 - \alpha(e))\hat{\mathbb{S}}(s_d)$ for some $\alpha(e) \in (0, 1)$, and for simplicity assume that $\alpha(e)$ is continuously differentiable (as are other relevant functions of e or $\mathbb{S}$ considered below).

Using this representation, a necessary and sufficient condition for $\hat{\mathbb{S}}(\hat{s})$ to be (strictly) monotonically increasing in e is that

$$\alpha'(e)\Big(e - \hat{\mathbb{S}}(s_d)\Big) + \alpha(e) > 0. \tag{A-4}$$

For $e > \hat{\mathbb{S}}(s_d)$, this requires that the weight on the estimate, $\alpha(e)$, not fall dramatically given small increases in e. For $e < \hat{\mathbb{S}}(s_d)$, the weight on the estimate must not fall dramatically given small decreases in e. Taken together, $\hat{\mathbb{S}}(\hat{s})$ will be monotonic in e as long as the weight on the estimate does not fall dramatically given small increases in $|e - \hat{\mathbb{S}}(s_d)|$ (i.e., as e moves further from the default $\hat{\mathbb{S}}(s_d)$), as stated in the text. Note that one simple sufficient condition for (A-4) is the constant-weighting case ($\alpha(e) = \alpha$), since in this case $\alpha'(e) = 0$ and the condition reduces to $\alpha(e) > 0$, which is guaranteed by Assumption 3.

Meanwhile, for $\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] - \mathbb{S}(s)$ to be (strictly) monotonically decreasing in $\mathbb{S}(s)$, we must have that

$$\frac{d\left(\mathbb{E}\left[\hat{\mathbb{S}}(\hat{s}) \,\middle|\, s\right]\right)}{d\mathbb{S}} - 1 < 0.$$

Fix a direction s_d, so that conditioning on s is equivalent to conditioning on $\mathbb{S}$. Since

$\mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s] = \int \Big((1-\alpha(e))\hat{\mathbb{S}}(s_d) + \alpha(e)e\Big)\, p(e|s)\, de$, the above condition requires

$$\begin{aligned} &\int \Big((1-\alpha(e))\hat{\mathbb{S}}(s_d) + \alpha(e)e\Big) \frac{\partial p(e|s)}{\partial \mathbb{S}}\, de < 1 \\ \iff &\int \Big((1-\alpha(e))\hat{\mathbb{S}}(s_d) + \alpha(e)e\Big) \frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)} p(e|s)\, de < 1 \\ \iff &\mathbb{E}\left[\Big((1-\alpha(e))\hat{\mathbb{S}}(s_d) + \alpha(e)e\Big) \frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)} \,\middle|\, s\right] < 1, \end{aligned}$$

or equivalently that

$$\text{(A-5)} \qquad \mathrm{Cov}_s\left(\alpha(e)(e - \hat{\mathbb{S}}(s_d)), \frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)}\right) < 1 - \mathbb{E}[\hat{\mathbb{S}}(\hat{s})|s]\mathbb{E}\left[\frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)} \,\middle|\, s\right],$$

where $\mathrm{Cov}_s(\cdot,\cdot)$ is the covariance conditional on s. Note further that

$$\mathbb{E}\left[\frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)} \,\middle|\, s\right] = \int \frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)} p(e|s)\, de = \int \frac{\partial p(e|s)}{\partial \mathbb{S}} de = 0,$$

since the density must integrate to 1 for all $\mathbb{S}$. The monotonicity condition in (A-5) can therefore be simplified to

$$\text{(A-6)} \qquad \mathrm{Cov}_s\left(\alpha(e)(e - \hat{\mathbb{S}}(s_d)), \frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)}\right) < 1.$$

By Assumption 1(b), $\frac{\frac{\partial p(e|s)}{\partial \mathbb{S}}}{p(e|s)}$ increases in e. Monotonicity in the degree of over-/underreaction in $\mathbb{S}$ therefore requires that $\alpha(e)$ not increase dramatically with e, as stated in the text, so that the covariance on the left side of (A-5) is less than 1. One can verify that this condition is again immediately satisfied in the constant-weighting case.

A.2. Additional Discussion for Section II.B

This appendix briefly discusses the mapping between the general environment in Section II.A and the log-normal environment in II.B. First, it is straightforward to verify that Assumptions 1 and 2 are satisfied in the log-normal environment. Assumption 3 is slightly more complex. This assumption requires that $\hat{\mathbb{S}}(s_d) < \hat{\mathbb{S}}(\hat{s}) < e$ when $e > \hat{\mathbb{S}}(s_d)$, and $\hat{\mathbb{S}}(s_d) > \hat{\mathbb{S}}(\hat{s}) > e$ when $e < \hat{\mathbb{S}}(s_d)$. Given the updating rule in (3), this requires $\exp(\sigma_{\mathbb{S}}^2/2) < \frac{e}{\hat{\mathbb{S}}(s_d)}$ when

$e > \hat{\mathbb{S}}(s_d)$, and $\exp(\sigma_e^2/2) < \frac{\hat{\mathbb{S}}(s_d)}{e}$ when $e < \hat{\mathbb{S}}(s_d)$.[3] In this case, the posterior is strictly between the prior and estimate. Alternatively, to guarantee that Assumption 3 holds for all e, one could drop the unbiasedness requirement of Assumption 1(a) (i.e., $\mathbb{E}[e|\mathbb{S}] = \mathbb{S}$, which is an unimportant normalization) and assume $\log e \sim \mathcal{N}(\log \mathbb{S}, \sigma_e^2)$, in which case Assumption 3 will always hold.

However, even if Assumption 3 is not guaranteed to hold in this log-normal setting, this is unimportant for our main results on over- and underreaction. This is demonstrated in equation (4), which shows that the conclusions in Proposition 1 apply regardless, and the person accordingly overreacts to weak signals and underreacts to strong signals (with resulting switching point $\mathbb{S}^*$ discussed in the text) in this log-normal environment. We accordingly do not focus on the conditions under which the primitive assumptions hold; what is important is that the main results continue to hold in this setting.

A.3. Proofs and Additional Discussion for Section II.C

Prior Distortions: Incorrect Priors, Uncertain Priors, and Base-Rate Neglect. In the case of an incorrect prior belief $\hat{\pi}_0$ discussed at the beginning of Section II.C, we can calculate the belief change $|\text{logit}(\hat{\pi}_1(s)) - \text{logit}(\hat{\pi}_0)|$ when $\hat{\pi}_0$ is observed. Perceived signal strength still follows the predictions in Proposition 1. Under the maintained assumption that belief changes are monotonic in perceived signal strength (see footnote 11 in the main text), the overreaction to weak signals and underreaction to strong signals in Proposition 1 will thus continue to be reflected in the belief change $|\text{logit}(\hat{\pi}_1(s)) - \text{logit}(\hat{\pi}_0)|$.

We now consider the case in which the correct prior is uncertain. We can model this by adding a pre-period $t = -1$, and we assume that the person entered this previous period with a prior $\hat{\pi}_{-1}$ known with certainty, then observed a signal s_0 (with known direction s_{d_0}) and used a strength estimate e_0 to form $\hat{\mathbb{S}}_0(\hat{s}_0)$ and $\hat{\pi}_0(\hat{s}_0)$ following Bayes' rule given distributions $p(\mathbb{S}_0|s_{d_0})$ and $p(e_0|s_{d_0}, \mathbb{S}_0)$.[4] This post-estimation prior is then the center of a non-degenerate distribution for the correct prior $\pi_0(s_0)$, representing a situation with uncertainty over this correct prior. The person then observes s_1 and updates to $\hat{\pi}_1(\hat{s}_1)$ as before (again following Bayes' rule), with s_1 independent of s_0 conditional on θ, and with e_0 and e_1 depending only

[3]These conditions will hold for most draws of e given reasonably small variances. More formally, these conditions are satisfied almost surely in a small-noise limit in which $\sigma_e^2/\sigma_{\mathbb{S}}^2$ is fixed while $\sigma_e^2, \sigma_{\mathbb{S}}^2 \to 0$. A similar limit is considered, for example, in Khaw, Li, and Woodford (2021, Section 4 and Appendix G).

[4]As in Section II.B, we continue to assume quasi-Bayesian updating. This allows us to formalize the statement in the text that $\hat{\pi}_0$ incorporates all uncertainty about past signals. In the more general case considered in Section II.A, the statement that the person overreacts to weak signals and underreacts to strong signals in period 1 is almost tautological: as long as the belief change continues to be monotonic in perceived signal strength, and perceived signal strength in period 1 is formed following Assumptions 1–3, then the results hold immediately.

on s_0 and s_1, respectively.[5]

With this setup, applying Bayes' rule twice, the posterior given $\hat{s}_0$ and $\hat{s}_1$ is

$$\begin{aligned}\text{logit}(\hat{\pi}_1) &= \text{logit}(\hat{\pi}_{-1}) + \log\left(\frac{p(\hat{s}_0|\theta=1)}{p(\hat{s}_0|\theta=0)}\right) + \log\left(\frac{p(\hat{s}_1|\hat{s}_0,\theta=1)}{p(\hat{s}_1|\hat{s}_0,\theta=0)}\right)\\ &= \text{logit}(\hat{\pi}_0(\hat{s}_0)) + \log\left(\frac{p(\hat{s}_1|\hat{s}_0,\theta=1)}{p(\hat{s}_1|\hat{s}_0,\theta=0)}\right).\end{aligned}$$

Note that $p(\hat{s}_1|\hat{s}_0,\theta) = p(\hat{s}_1|\theta)$, since s_0 and s_1 are independent conditional on θ, and e_0 and e_1 depend only on s_0 and s_1, respectively. Therefore,

$$\text{logit}(\hat{\pi}_1) = \text{logit}(\hat{\pi}_0(\hat{s}_0)) + \log\left(\frac{p(\hat{s}_1|\theta=1)}{p(\hat{s}_1|\theta=0)}\right).$$

The belief update in period 1, $|\text{logit}(\hat{\pi}_1) - \text{logit}(\hat{\pi}_0(\hat{s}_0))|$, accordingly depends on perceived signal strength $\left|\log\left(\frac{p(\hat{s}_1|\theta=1)}{p(\hat{s}_1|\theta=0)}\right)\right|$ exactly as was the case before, with the previous period's estimate (or multiple previous periods' estimates) affecting only $\hat{\pi}_0$. Under the assumption that $\left|\log\left(\frac{p(\hat{s}_1|\theta=1)}{p(\hat{s}_1|\theta=0)}\right)\right|$ is monotonic in $\hat{\mathbb{S}}_1(\hat{s}_1) = \hat{\mathbb{E}}[\mathbb{S}_1|s_{d_1}, e_1]$ (again as in main text footnote 11), all our results therefore carry through to this case.

In the case that the previously formed prior is unobserved, though, we cannot calculate $|\text{logit}(\hat{\pi}_1) - \text{logit}(\hat{\pi}_0(\hat{s}_0))|$ directly. Instead, we again use $|\text{logit}(\hat{\pi}_1) - \text{logit}(\pi_0)|$ as our proxy for reaction. This measure now includes both perceived signal strength and the prior distortion:

$$\text{(A-7)} \qquad |\text{logit}(\hat{\pi}_1) - \text{logit}(\pi_0)| = |\text{logit}(\hat{\pi}_0(\hat{s}_0)) - \text{logit}(\pi_0)| \pm \left|\log\left(\frac{p(\hat{s}_1|\theta=1)}{p(\hat{s}_1|\theta=0)}\right)\right|.$$

There are thus two cases to consider. (1) If the expected prior distortion in the first term has the same sign as the signal direction, then $|\text{logit}(\hat{\pi}_1) - \text{logit}(\pi_0)|$ will overstate the degree of overreaction in the perceived signal strength $\hat{\mathbb{S}}(\hat{s})$, and there may appear to be overreaction even to strong signals. This will apply, for example, if the correct prior is much lower than 0.5, but people do not use this correct prior and instead shade toward a default uninformative prior of 0.5. This will push up the apparent reaction to a positive signal. (2) If the expected prior distortion has the opposite sign as the signal direction, then $|\text{logit}(\hat{\pi}_1) - \text{logit}(\pi_0)|$ will understate the degree of overreaction in the perceived signal strength $\hat{\mathbb{S}}(\hat{s})$, and there may appear to be underreaction (or incorrectly signed reactions) even to weak signals. Intuitively, the prior distortion offsets the signal reaction in this case. We should expect these issues to

[5]Note that this setup does not depend on the specific timing of periods 0 and 1; this notation simply formalizes the idea that the correct prior is formed from some signal (like information provided in an experiment) separate from the additional piece of information in signal s_1.

matter less when the prior estimation is more precise than the signal strength estimation, or when the default prior (often 0.5) is close to the correct prior.

The same analysis applies to the case with base-rate neglect, which will simply move the effective prior $\hat{\pi}_0$ in (A-7) toward the person's default prior (which, in this binary-state setting, is again often modeled as the uninformative prior of 0.5). Cases in which the correct prior is equal to or close to 0.5 will therefore have little to no role for such base-rate neglect. More generally, we expect our results to hold within a range of priors around $\pi_0 = 0.5$. (Based on our experimental results, this range appears reasonably wide.) For correct priors close to 0 or 1, meanwhile, given strong enough base-rate neglect, this can offset our main effect according to situations (1) and (2) as described in the preceding paragraph.

We discuss and control for the effects of base-rate neglect in additional detail in Section III.B, which presents the results of an experiment with priors different from 0.5. In particular, as shown in eq. (8) and discussed in footnote 28 of the main text, the measured signal weight $\hat{w}(s)$ (which we estimate as $\frac{\operatorname{logit}\hat{\pi}_1 - \operatorname{logit}\pi_0}{\operatorname{logit}\pi_1 - \operatorname{logit}\pi_0}$) will be distorted by base-rate neglect to the extent that $\operatorname{logit}\pi_0$ (the distance of the prior from 0.5) is high relative to $\operatorname{logit}\pi_1 - \operatorname{logit}\pi_0$ (the true signed signal strength), though of course this only matters to the degree that the person engages in strong base-rate neglect.

Uncertainty About the Direction. Following the discussion in the text, we now assume that the person forms an estimate e of $\mathbb{S}_{signed}$, with that estimate satisfying Assumption 1 with respect to $\mathbb{S}_{signed}$. In place of Assumption 2, we assume that the default value (the person's subjective prior) is $\hat{\mathbb{S}}_{0,signed} = 0$. This effectively assumes a symmetric signal strength distribution where $\mathbb{E}[\mathbb{S}_{signed}] = 0$.[6] Similar to Assumption 3, we assume that the posterior $\hat{\mathbb{S}}_{signed}(\hat{s})$ is strictly between 0 and the estimate e. Given this, it is immediate that on average, there is underreaction in perceived signed strength: $\mathbb{E}[|\hat{\mathbb{S}}_{signed}| \mid s] = a(s)|\mathbb{S}_{signed}| < |\mathbb{S}_{signed}|$, where $a(s) \in (0,1)$. Note that this definition of underreaction is in terms of the absolute perceived signed strength relative to the absolute true signed strength. The interpretation of this result as "underreaction" becomes more strained when the signs of $\hat{\mathbb{S}}_{signed}$ and $\mathbb{S}_{signed}$ are different, as discussed in the main text.

[6] We note that the prediction of underreaction does not necessarily apply in asymmetric cases, which can lead to strange situations. While the expected change in beliefs is always equal to 0 (at least for a Bayesian), due to the non-linear transformation from signal strength to belief changes, it is not necessarily the case that the signal strength has a mean of zero. If one removes the assumption of symmetry, we can say only that there is underreaction for sufficiently extreme signed strengths, but we cannot necessarily make statements across all signal strengths. This analysis is, however, not the main focus given the settings we seek to describe.

A.4. Proofs and Additional Discussion for Section II.D

Independent Estimates. We formally define the cross-sectional expectation as $\mathbb{E}_i[X_i|s] = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} X_i$ for any measurable X_i (whose distribution implicitly depends on s). Under the assumptions that the estimates e_i are independent across people, there is no formal distinction between taking the expectation with respect to the distribution of estimates (as we did previously) and taking the cross-sectional expectation across people. Therefore, Proposition 1 continues to apply, in the sense that there exists a unique switching point $\mathbb{S}^*$ such that there is overreaction on average ($\mathbb{E}_i[\hat{\mathbb{S}}_i(\hat{s}_i)|s] > \mathbb{S}(s)$) if $\mathbb{S}(s) < \mathbb{S}^*$, and there is underreaction on average ($\mathbb{E}_i[\hat{\mathbb{S}}_i(\hat{s}_i)|s] < \mathbb{S}(s)$) if $\mathbb{S}(s) > \mathbb{S}^*$.

Correlated Estimates and Over-/Underreaction Conditional on $\mathbb{S}$*.* We consider the case with multi-dimensional signals and perfect correlation in estimates given identical attention vectors $\boldsymbol{a}_i$ for all i. In this case, as stated in the text, Proposition 1 holds under the following new definition: there is overreaction if $\mathbb{E}[\hat{\mathbb{S}}_i(\hat{s}_i)|\mathbb{S}] > \mathbb{S}$, and underreaction if $\mathbb{E}[\hat{\mathbb{S}}_i(\hat{s}_i)|\mathbb{S}] < \mathbb{S}$.

Since $s_{m,j}$ are i.i.d. over components j and exchangeable, $\mathbb{E}[s_{m,j}|\mathbb{S}] = \log \mathbb{S}$. (By comparison, conditional on s, $s_{m,j}$ is known, so a given person's $\hat{\mathbb{S}}(s)$ in that case could potentially differ across s for the same $\mathbb{S}$, invalidating our results. This motivates our conditioning on $\mathbb{S}$ here.) Similarly, e_i is log-normally distributed conditional on $\mathbb{S}$, $\log e_i \sim \mathcal{N}(\log \mathbb{S} - \sigma_{e,i}^2/2, \sigma_{e,i}^2)$, where this (and the expression for $\sigma_{e,i}^2$ provided in the text) follow from standard characterizations of a multivariate normal distribution along with some algebra. Therefore, conditioning on $\mathbb{S}$, $\mathbb{E}[\hat{\mathbb{S}}_i(\hat{s}_i)|\mathbb{S}] = k\mathbb{S}^\beta$, where k and β are as given in the text. This is exactly as in (4), and we conclude that Proposition 1 applies using the above definition of over- and underreaction. Further, since we have assumed the extreme case of perfectly correlated estimates, this will also apply when considering the expected cross-sectional expectation $\mathbb{E}[\mathbb{E}_i[\hat{\mathbb{S}}_i(\hat{s}_i)|s]|\mathbb{S}]$ given that $\hat{\mathbb{S}}_i(\hat{s}_i)$ is identical across i for a given s.

Predictions on Correlation Behavior. As stated in the text, one can make further statements about the correlation in updating behavior across people under additional assumptions about the signal components and person-specific vectors $\boldsymbol{a}_i$. For example, if the components are ordered by salience, then it is natural to assume that $\boldsymbol{a}_i$ is such that $a_{i,j} = 1$ for $j \le n_i$ and $a_{i,j} = 0$ for $n_i < j \le n$ (i.e., person i pays attention to the first n_i components, and the only difference across people is in how large n_i is). In this case, the following expression holds for the ex ante correlation between estimates for any two people i and i', ordered such that $0 < n_i \le n_{i'}$:

$$\mathrm{Corr}(e_i, e_{i'}) = \sqrt{\frac{n_i}{n_{i'}}} \in (0, 1].$$

This expression follows from the fact that $\text{Cov}(e_i, e_{i'}) = \text{Var}(s_{m,j})/\max(n_i, n_{i'})$, while $\text{Var}(e_i) = \text{Var}(s_{m,j})/n_i$ and $\text{Var}(e_{i'}) = \text{Var}(s_{m,j})/n_{i'}$.

The above expression can be generalized to cases where the components are not salience-ordered. In these cases, the correlation will simply scale down as one decreases the overlap in the entries of $\boldsymbol{a}_i$ and $\boldsymbol{a}_{i'}$ that are equal to 1. For example, if there are two components and two types of people, with type 1 attentive only to component 1 and type 2 attentive only to component 2, then there will be multimodal estimates, with perfect correlation across people within type and none across types. With high-dimensional vectors in which the components are not ordered according to salience (i.e., cases where people's attention vectors are varied), estimates will be closer to the independent case, and we will see smoother distributions of resulting strength perceptions.

A.5. Proofs for Section IV.A

Proof of Proposition 2. Here, we provide a brief restatement of the proof of Proposition 1 of AR (2021) for completeness. Since $\pi_t = \pi_t(H_t) = \mathbb{E}_t[\theta]$, by the law of iterated expectations (LIE), beliefs are a martingale: $\pi_t = \mathbb{E}_t[\pi_{t+1}]$. Therefore, for arbitrary t_1,

$$\begin{aligned}
\mathbb{E}_{t_1}[M_{t_1,t_1+1} - R_{t_1,t_1+1}] &= \mathbb{E}_{t_1}[(\pi_{t_1+1} - \pi_{t_1})^2 - (\pi_{t_1}(1-\pi_{t_1}) - \pi_{t_1+1}(1-\pi_{t_1+1}))] \\
&= \mathbb{E}_{t_1}[(2\pi_{t_1} - 1)(\pi_{t_1} - \pi_{t_1+1})] \\
&= (2\pi_{t_1} - 1)(\mathbb{E}_{t_1}[\pi_{t_1} - \pi_{t_1+1}]) = 0,
\end{aligned}$$

where the first line uses the definition of movement and uncertainty reduction, the second line simplifies, and the last line rearranges and uses the martingale property of beliefs. Similarly, $\mathbb{E}_{t_1+\tau}[M_{t_1+\tau,t_1+\tau+1} - R_{t_1+\tau,t_1+\tau+1}] = 0$ for any $\tau \geq 0$, and therefore by the LIE, $\mathbb{E}_{t_1}[M_{t_1+\tau,t_1+\tau+1} - R_{t_1+\tau,t_1+\tau+1}] = 0$. So summing all these terms from t_1 to arbitrary $t_2 > t_1$,

$$\mathbb{E}_{t_1}[M_{t_1,t_2} - R_{t_1,t_2}] = \sum_{\tau=0}^{t_2-t_1-1} \mathbb{E}_{t_1}[M_{t_1+\tau,t_1+\tau+1} - R_{t_1+\tau,t_1+\tau+1}] = 0,$$

as stated. $\square$

B. Additional Data Details and Estimation Results

B.1. Experimental Studies: Details and Robustness Checks

Study 1a

Timing Details. Participants saw the following five treatment blocks: (1) one symmetric signal, (2) one asymmetric signal, (3) three symmetric signals, (4) demand for information, (5) uncertain signals. Details of each are in the subsequent subsections. The ordering of when they saw each treatment block was as follows:

Rounds	**Treatment Block**	**Frequency**	**Observations**
1–12	One symmetric signal	67 percent	4,036
1–12	One asymmetric signal	33 percent	1,964
13	Attention check	100 percent	500
14–18	Three symmetric signals	100 percent	2,500
19–23	Demand for information	100 percent	2,500
24–25	One uncertain signal	100 percent	1,000

Questions within each treatment block were randomized for each participant. The ordering of treatment blocks (besides "one symmetric" and "one asymmetric") were fixed for ease of participant comprehension. For instance, participants do not see the "demand for information" treatment until they have played rounds in which they inferred from one signal and from multiple signals. The uncertain-signals treatment comes after the demand-for-information treatment because they do not reflect the signals that participants would purchase.

Additional Results Discussed in the Text. Figures A1–A3 provide additional results discussed in the text.

FIGURE A1. Comparison to Literature

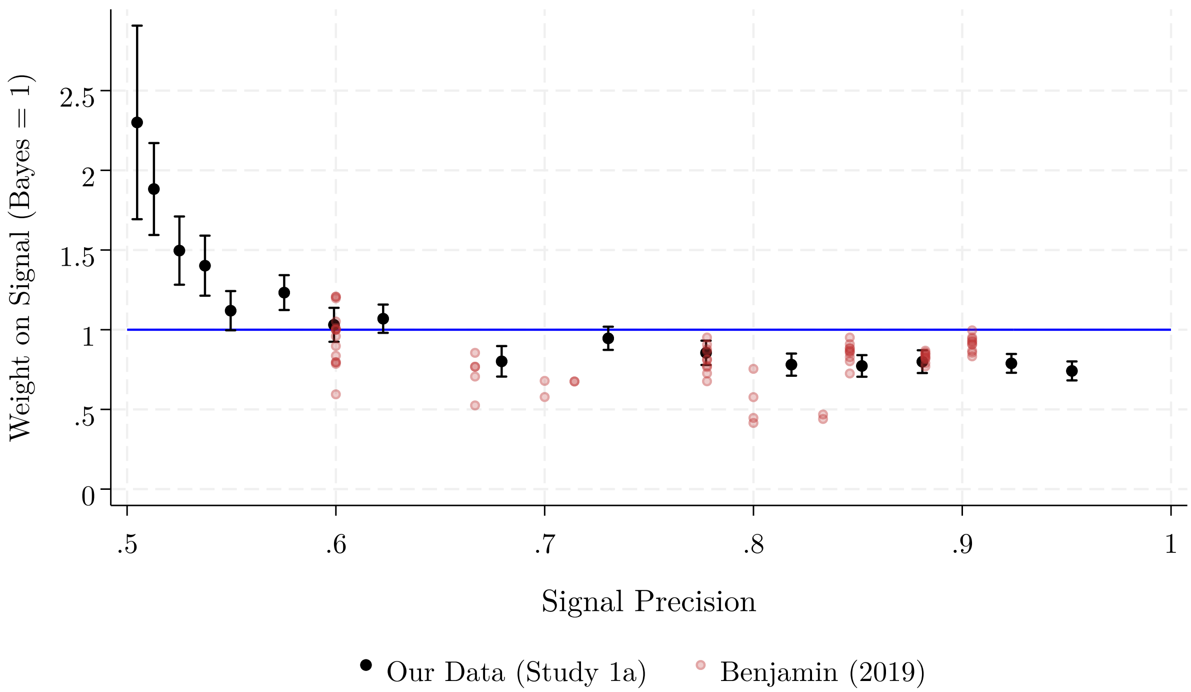

Notes: This figure shows the weight put on signals of different precisions, where weight is defined relative to a Bayesian (whose weight of 1 is in the blue line) as in the main text. Black circles correspond to data from our Study 1a with 95% confidence intervals (as plotted in Figure II). Light translucent circles correspond to data from Benjamin (2019). We use the data from his supplementary files, restricting to the 70 studies in which participants update from one binary signal when the prior is 0.5 and signal precision is symmetric. Note that most papers include multiple studies.

Figure A2. Heterogeneity in Inference at the Individual Level

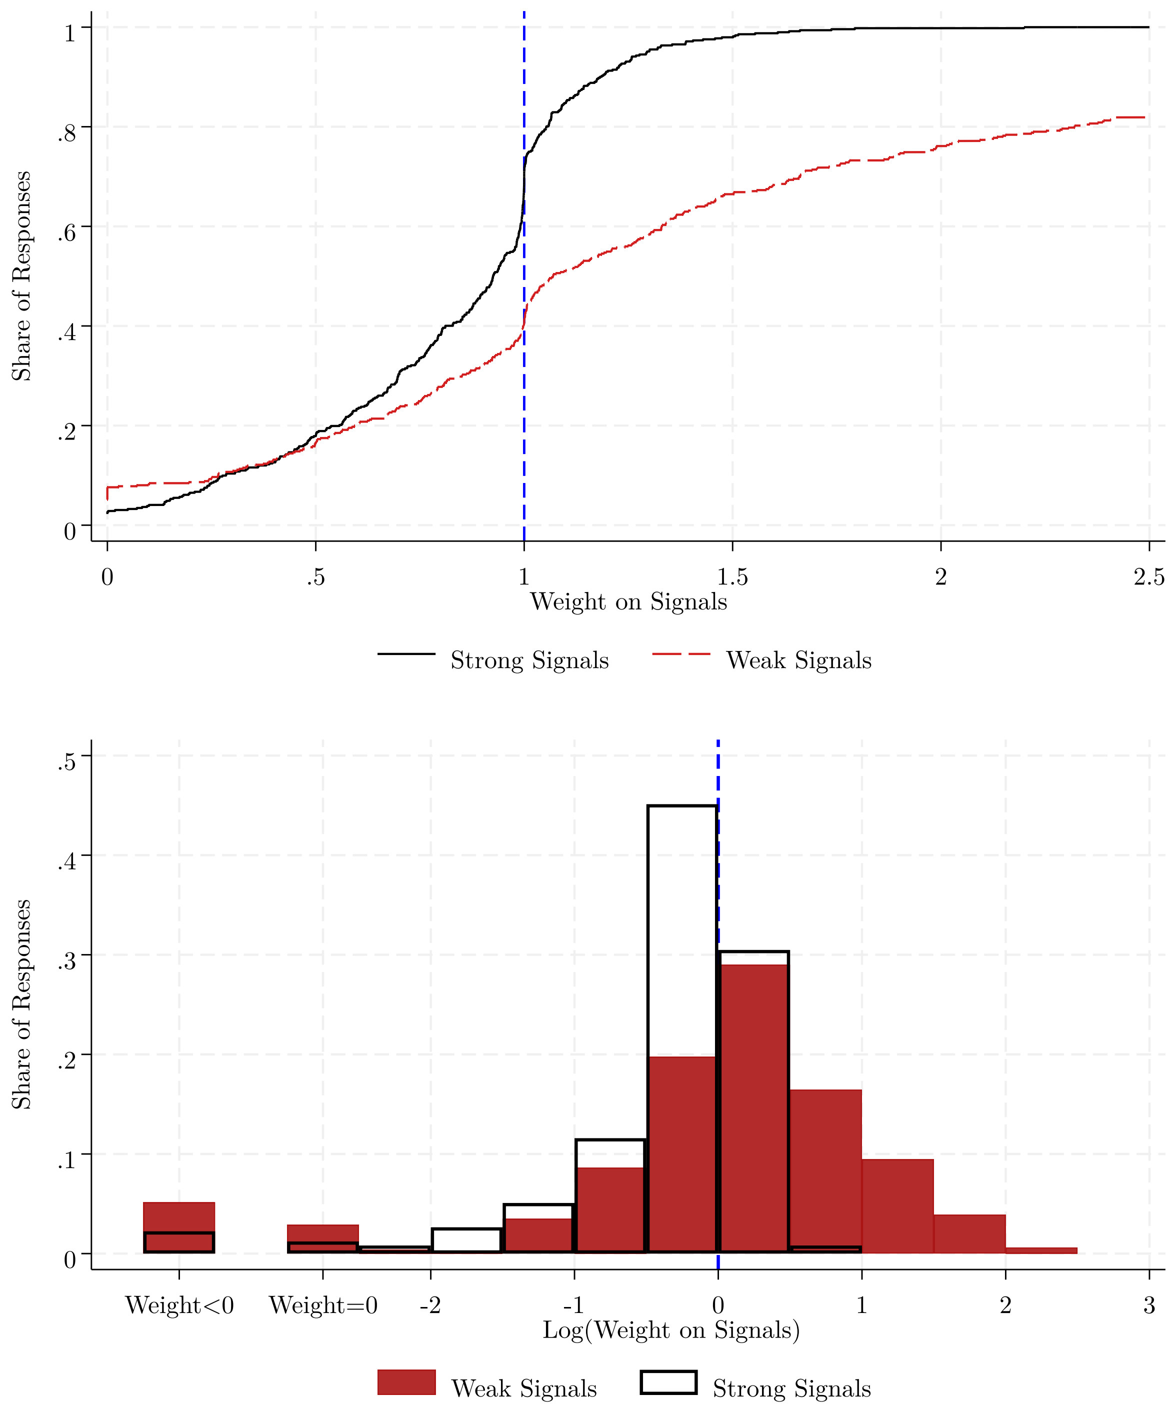

Notes: This figure shows how much weight is put on weak and strong signals at the individual level, where weight is defined relative to a Bayesian as in the main text. The top panel shows the CDF of individuals' weights on strong and weak signals. The vertical line at 1 represents Bayesian updating. The bottom panel shows the PDF of the log of individuals' weights on strong and weak signals. The vertical line at 0 represents Bayesian updating. Participants with nonpositive weight are separated out. Weak signals have precision $p < 0.6$ and strong signals have precision $p > 0.7$. Observations are winsorized, for each signal strength, at the 5% level.

Figure A3. Over- and Underinference by Number and Strength of Signals

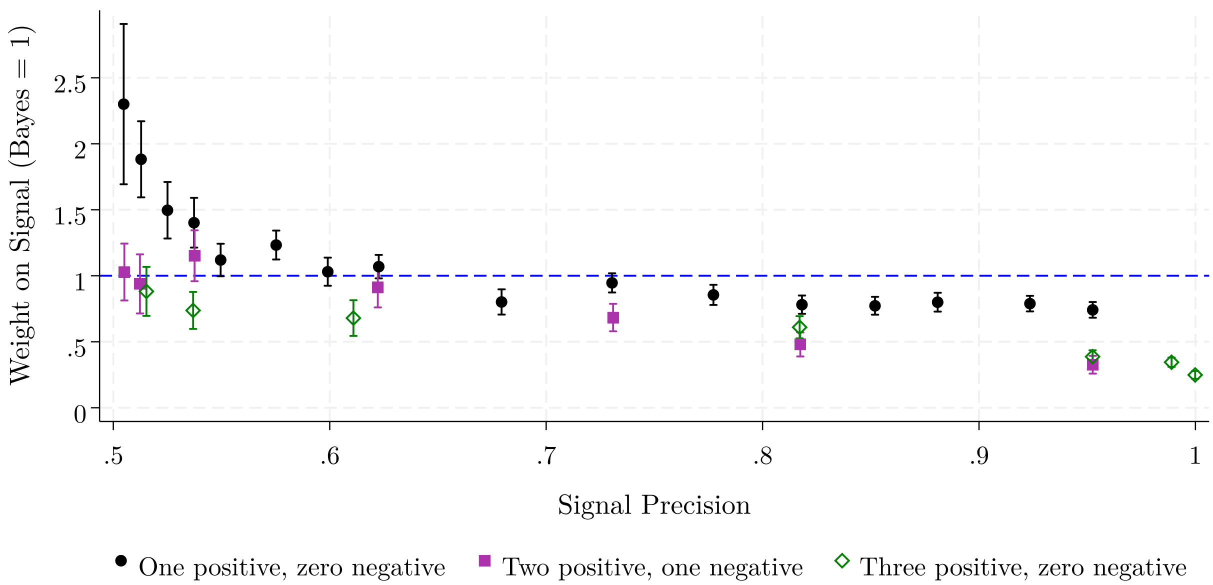

Notes: This figure plots the average weight participants put on signals relative to a Bayesian (indicated by the dashed line), split by signal distribution. Black circles correspond to one signal of precision ρ (as in Figure II); light squares correspond to two signals of precision ρ in one direction and one signal of precision ρ in the opposing direction; and hollow diamonds correspond to three signals of strength $\mathbb{S}/3$, where $\mathbb{S} = \text{logit}\rho$ for precision ρ, in the same the direction. This figure shows that participants put less weight on three signals as compared to the weight they put on one signal but that weight declines in signal precision in all cases. Error bars indicate 95% confidence intervals.

Further Results: Demand for Information. Patterns of overinference and underinference can also lead to demand for information that is too high or too low relative to the optimum. Figure A4 plots the average number of signals purchased as a function of each signal strength, comparing participant behavior to the optimal choice if participants were Bayesian and only valued signals for their instrumental value.

FIGURE A4. Number of Signals Purchased

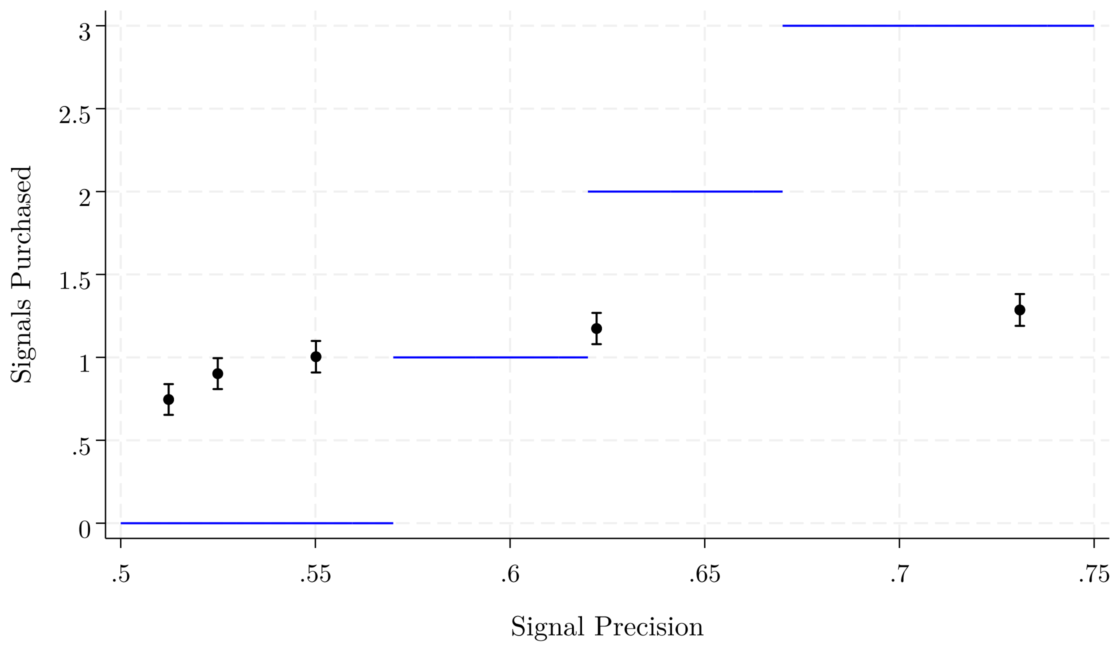

Notes: This figure plots the number of signals purchased as a function of signal precision. The horizontal lines correspond to the payoff-maximizing number of signals that would be purchased. This figure shows that participants over-purchase weak signals and under-purchase strong signals relative to a payoff maximizer. Error bars indicate 95% confidence intervals.

As can be seen in the figure, participants systematically over-purchase weak signals and under-purchase strong signals. The cost of a signal that leads a Bayesian to form a posterior of less than 0.57 outweighs its benefit; however, the majority of participants purchase at least one signal when $p = 0.55$ and $p = 0.525$. Additionally, 81 percent of participants purchase fewer than the optimal level of three signals when $p = 0.73$. Over- and underinference patterns therefore matter not just for stated beliefs; they also lead people to overvalue low-quality information and undervalue high-quality information, as reflected in their purchase decisions.

Study 1b

Grether Regression Approach

TABLE A1. Effect of Logit Prior and Signal Strength on Logit Posterior

	(1) 50% Prior	(2) All Priors	(3) All Priors	(4) All Priors
Signal Strength: 0.05	2.371	2.844	2.800	2.795
	(0.224)	(0.286)	(0.179)	(0.178)
Signal Strength: 0.20	1.359	1.554	1.504	1.500
	(0.072)	(0.083)	(0.059)	(0.059)
Signal Strength: 0.50	1.176	1.208	1.190	1.191
	(0.054)	(0.041)	(0.035)	(0.035)
Signal Strength: 1.25	0.840	0.852	0.857	0.857
	(0.028)	(0.021)	(0.020)	(0.020)
Signal Strength: 1.75	0.824	0.768	0.762	0.762
	(0.021)	(0.017)	(0.015)	(0.015)
Logit Prior		1	0.984	
		(.)	(0.016)	
Logit Prior: Signal=0.05				1.028
				(0.017)
Logit Prior: Signal=0.20				1.036
				(0.021)
Logit Prior: Signal=0.50				0.996
				(0.023)
Logit Prior: Signal=1.25				0.944
				(0.029)
Logit Prior: Signal=1.75				0.915
				(0.029)
Participant FE	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes
Observations	2500	7500	7500	7500
R^2	0.80	0.60	0.76	0.76

Notes: OLS, with standard errors in parentheses clustered at participant level. We regress logit posterior on each signal strength separately. Column (1) restricts to observations where the prior is symmetric (as in Study 1a); other columns use the full dataset. Column (2) assumes that people put weight 1 on their prior. Column (3) allows for misweighting priors overall. Column (4) allows for weights on priors to vary for each signal strength. See main text for discussion.

Study 2

Grether Regression Approach

TABLE A2. Weight on Signal and Prior by Quarter of Basketball Game

	(1) All Quarters	(2) All Quarters	(3) All Quarters
Quarter 1 x Signal Strength	1.344	1.405	1.406
	(0.108)	(0.066)	(0.066)
Quarter 2 x Signal Strength	1.398	1.360	1.359
	(0.110)	(0.059)	(0.059)
Quarter 3 x Signal Strength	0.968	0.929	0.928
	(0.070)	(0.041)	(0.040)
Quarter 4 x Signal Strength	0.735	0.585	0.587
	(0.037)	(0.020)	(0.020)
Logit Prior	1	0.906	
	(.)	(0.013)	
Quarter 1 x Logit Prior			1.001
			(0.051)
Quarter 2 x Logit Prior			0.948
			(0.031)
Quarter 3 x Logit Prior			0.917
			(0.025)
Quarter 4 x Logit Prior			0.889
			(0.014)
Participant FE	Yes	Yes	Yes
Round FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	8000	8000	8000
R^2	0.48	0.86	0.86

Notes: OLS, with standard errors in parentheses clustered at participant level. We regress logit posterior on signals in each quarter separately. Column (1) assumes that people put weight 1 on their prior. Column (2) allows for misweighting priors overall. Column (3) allows for weights on priors to vary for each quarter. See main text for discussion.

B.2. Empirical Analysis: Details and Robustness Checks

Measurement Details for Risk-Neutral Beliefs

This subsection describes our use of option-price data, as introduced in Section IV.C, in greater detail. (Much of this detail is directly from AL 2023.) First, we describe how we clean the option data and then translate the option prices to risk-neutral beliefs. We then detail how we translate from risk-neutral to physical beliefs under different parameterizations for risk aversion.

Option Data Cleaning and the Risk-Neutral Distribution. We start from the OptionMetrics data described in the text, obtaining the end-of-day bid and ask prices for all European call and put options on the S&P 500 index, for all available strike prices and option expiration dates for trading dates from January 1996 through December 2018. We then average the bid and ask price to obtain the mid price. We also, as in AL (2023), obtain S&P 500 index prices to use when determining the realized index-return state. We first get end-of-day index prices (which we take as well from OptionMetrics, and then augment these with hand-collected settlement values for any options whose settlement value depends on the opening (rather than closing) index price, from the CBOE website.[7]

To measure the risk-free rate $R^f_{t,T}$ in order to define our excess return space, we follow van Binsbergen, Diamond, and Grotteria (2022) and estimate the risk-free rate from the cross-section of option prices by applying put-call parity. We use their "Estimator 2," which estimates $R^f_{t,T}$ from Theil–Sen (robust median) estimation of the put-call parity relationship. This provides a risk-free rate consistent with observed option prices.

For the OptionMetrics data, we then use the same steps as described in Online Appendix C.5 of AL (2023) to clean the data and convert to a risk-neutral distribution. For cleaning, we drop any options with bid or ask price of zero (or less than zero), with uncomputable Black–Scholes implied volatility or with implied volatility of greater than 100 percent, with more than one year to maturity, or (for call options) with mid prices greater than the price of the underlying; we drop any option cross-section (i.e., the full set of prices for the pair (t, T)) with no trading volume on date t, with fewer than three listed prices across different strikes, or for which there are fewer than three strikes for which both call and put prices are available (as is necessary to calculate the forward price and risk-free rate).

We then measure the risk-neutral distribution following Malz (2014), again as described

[7]The results for the binarized noise-corrected data in Figure A10 below also use separate data directly from AL (2023), so we refer to that paper — in particular, Section 6 and Online Appendix C.5–C.6 — for details on the data and methodology used for the noise estimation (which use intraday option data obtained directly from the CBOE), as well as the conversion of the histogram of risk-neutral beliefs to binarized beliefs.

in Online Appendix C.5 of AL (2023):

1. We translate the option mid prices into equivalent Black–Scholes implied volatilities.
2. We discard the resulting observations for in-the-money calls and puts, so that the remaining steps use data from only out-of-the-money put and call prices. To determine the at-the-money point, we use the strike K at which call and put prices are equal (or closest to each other).
3. For each trading date–expiration date pair, we fit a clamped cubic spline to the resulting implied volatility curve (i.e., the curve of implied volatility vs. strike price).
4. Evaluate this spline at 1,901 strike prices, for S&P index values ranging from 200 to 4,000 (so that the evaluation strike prices are $K = 200, 202, \ldots, 4000$), to obtain a set of fitted implied-volatility values across this fine grid of possible strike prices for each (t, T) pair.
5. Invert the resulting smoothed implied volatility schedule back into call prices $\hat{q}_{t,T,K}$.
6. Using a discrete-state version of the classic Breeden and Litzenberger (1978) formula, calculate the risk-neutral CDF for the date-T index value at strike price K as follows: $\mathbb{P}_t^*(V_T < K) = 1 + R_{t,T}^f(\hat{q}_{t,T,K} - \hat{q}_{t,T,K-2})/2$.
7. Defining $V_{T,j,\max}$ and $V_{T,j,\min}$ to be the date-T index values corresponding to the upper and lower bounds, respectively, of the bin defining return state θ_j (i.e., the upper and lower end of the five-percentage-point excess-return range defining a given return outcome), calculate the risk-neutral belief that state θ_j will be realized at date T as $\pi_{t,j}^* = \mathbb{P}_t^*(V_T < V_{T,j,\max}) - \mathbb{P}_t^*(V_T < V_{T,j,\min})$. (The beliefs for states θ_1 and θ_{10} then collect the tail probabilities for below -20% and above 20% returns, respectively.)

We do this for states $\theta_1, \ldots, \theta_{10}$, where the return states are as defined and described in the text — i.e., 5-log-point ranges of log excess returns from the first observable option trading date (within a year of expiration) to the expiration date — for all trading dates under consideration. We then use the resulting histogram of risk-neutral beliefs for our tests.

We note that unlike AL (2023), we include beliefs over the tail return states θ_1 and θ_{10}, whereas AL discard them before calculating binarized beliefs $\pi_{t,j}^*/(\pi_{t,j}^* + \pi_{t,j+1}^*)$. AL discard them due to concerns over complications from potential changes in risk aversion over tail outcomes; given the binarization, small changes in risk aversion would have large effects on the measured binarized RN beliefs. But this is not the case for our analysis: since we just use the (non-binarized) histogram of beliefs, the tail states have very low probabilities and thus do not meaningfully affect the results. (Results are very similar when only including movement and uncertainty reduction for states 2 through 9.) This is another way in which just using the RN histogram, rather than continuing from above and calculating the binarized beliefs, helps minimize the potential effect of noise on our results.

Translating from Risk-Neutral to Physical Beliefs. Given the RN beliefs as measured from above, we now describe the translation from RN to physical beliefs in greater detail. Assume there exists a representative investor ("the market") with time-separable utility over the market index value.[8] Assume, as above, that the state space (the set of possible terminal index values V_T) is discrete, with states indexed by j ($V_T = \theta_j$ for $j = 1, 2, \ldots, J$), and denote terminal utility by $U(V_T)$. The physical belief regarding the likelihood of state j is $\pi_{t,j}$, and the risk-neutral belief is $\pi^*_{t,j}$. The two are related as follows:

$$\pi^*_{t,j} = \frac{U'(\theta_j)\pi_{t,j}}{\sum_k U'(\theta_k)\pi_{t,k}}. \tag{A-8}$$

(This is a multi-state generalization of equation (5) of AL 2023, or see equation (7) of Bliss and Panigirtzoglou 2004.) Our main translation assumes that $U'(V_T) = V_T^{-\gamma}$, corresponding to the assumption of power utility over the terminal index return, with constant relative risk aversion coefficient of γ. We then follow Bliss and Panigirtzoglou (2004) in estimating γ as the value under which the physical beliefs over the S&P 500 value at the one-month horizon are well calibrated (i.e., unbiased); see Bliss and Panigirtzoglou (2004) for details on the maximum likelihood estimation procedure.[9]

We then consider dozens of generalizations of this basic framework. First, we reparameterize (A-8) in terms of the *ratio* of marginal utilities (or SDF realizations) across adjacent index states ϕ_j, by substituting $U'(\theta_j) = \phi_j U'(\theta_{j-1})$. We then make a range of assumptions on the function ϕ_j. We assume that ϕ_j varies by state j, either linearly or quadratically in V_T, and we estimate ϕ_j by maximum likelihood for each state; we assume that ϕ_j varies over time (either linearly or quadratically) or by horizon to expiration (as in Lazarus 2022); and then we consider interactions in which ϕ_j varies both by bin j and over time. In all cases (as can be seen in Figure A9, the right panel of which contains one line for each parameterization), the movement and uncertainty reduction statistics are close to unchanged. (This is in contrast to the *physical probabilities*, which do change depending on the parameterization; it is their evolution over time that is unchanged.)

[8] These illustrative assumptions aid in the interpretation of our risk-aversion assumptions, but they are stronger than needed in general; see AL (2023) for a discussion.

[9] Like Bliss and Panigirtzoglou (2004), we obtain reasonable estimates of risk aversion of (with estimated $\hat{\gamma} < 10$) given this calibration procedure.

Additional Empirical Results

We now provide a set of robustness results (Figures A5–A12 and Tables A3–A4) for the betting and finance data. As described in the text, we present figures and regression tables for when the data is split into either 12 or 36 time chunks. For the options data, we also show results of different risk adjustments, use of binarized noise-corrected data, and in subsamples.

Different Time Windows

Figure A5. Movement and Uncertainty Reduction for Sports Betting: 12 Time Chunks

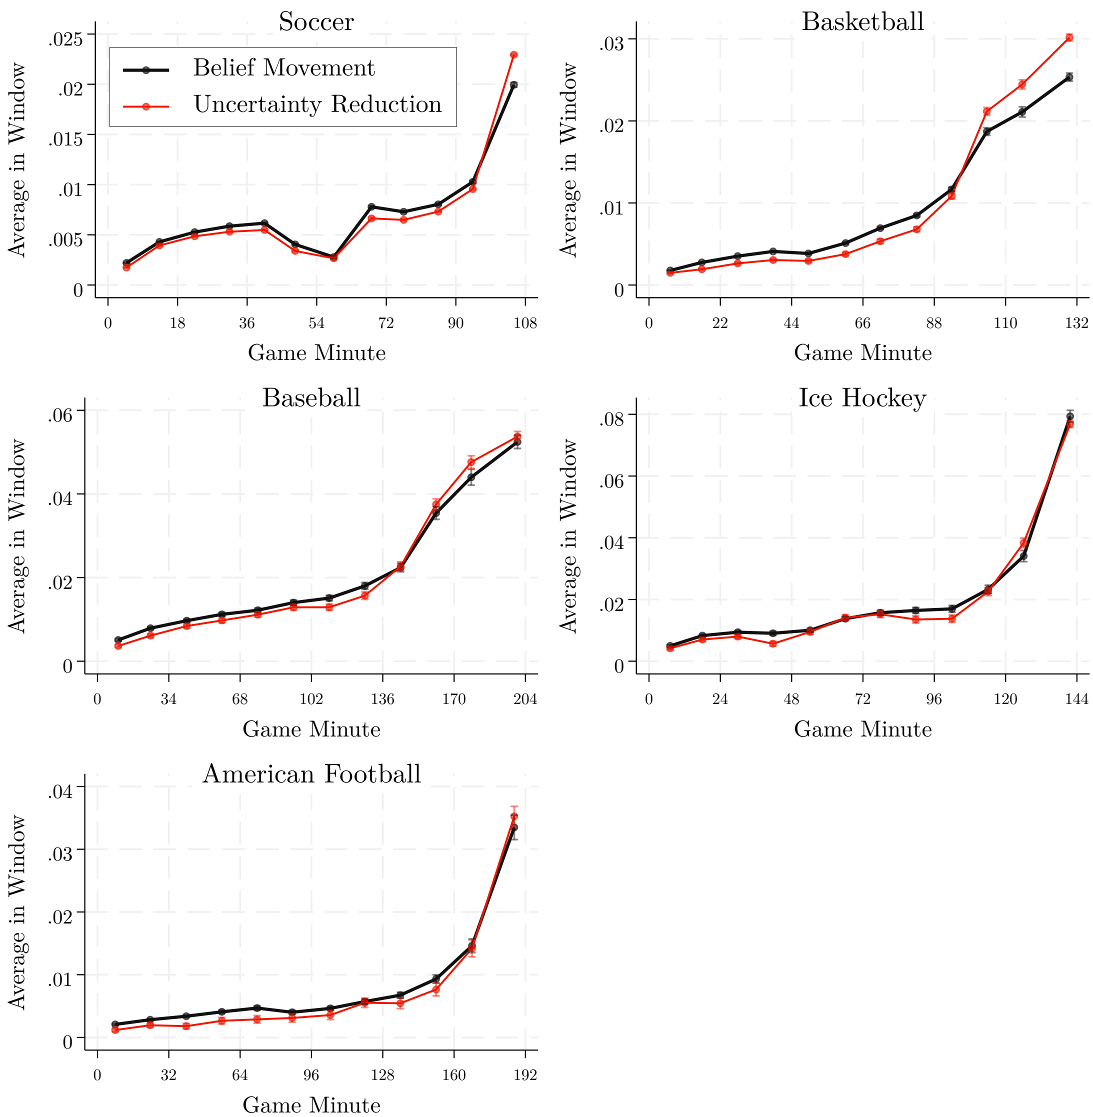

Notes: This figure replicates Figure VII, but with 12 equal-length time windows, rather than 24. See that figure's notes for details on construction.

Figure A6. Movement and Uncertainty Reduction for Finance Data: 12 Time Chunks

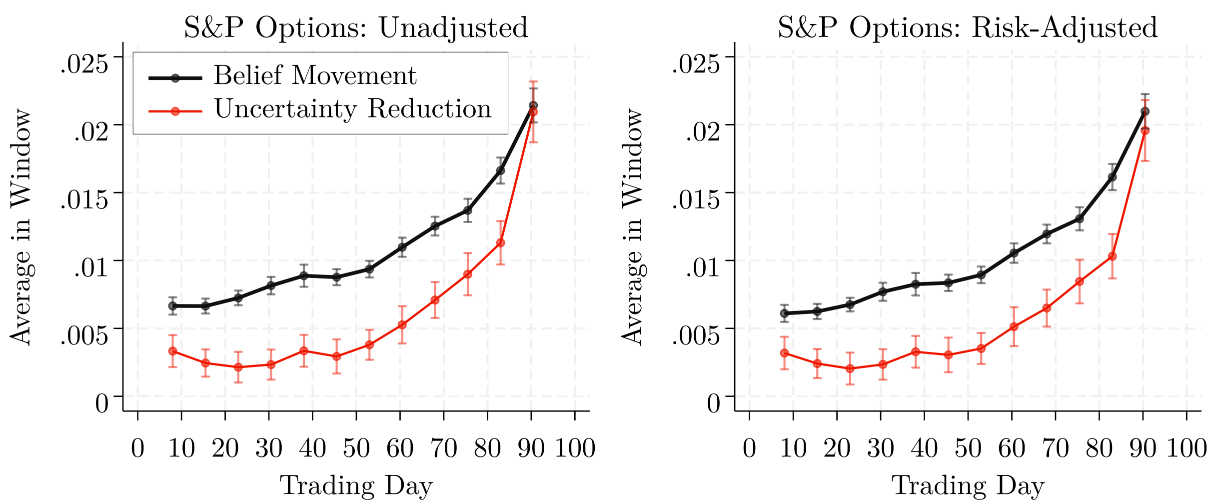

Notes: This figure replicates Figure VIII, but with 12 equal-length time windows, rather than 24. See that figure's notes for details on construction.

Table A3. Regressions of Movement on Uncertainty Reduction: 12 Time Chunks

Dep. Var.: Movement	Sports					Finance	
	Soccer	Basketball	Baseball	Hockey	Football	Raw	Risk-Adj.
Constant	0.0014	0.0018	0.0024	0.0013	0.0015	0.0060	0.0054
	(0.0003)	(0.0003)	(0.0004)	(0.0009)	(0.0002)	(0.0005)	(0.0005)
Uncert. Red.	0.839	0.797	0.903	0.987	0.912	0.796	0.861
	(0.006)	(0.007)	(0.012)	(0.012)	(0.027)	(0.054)	(0.063)
R^2	0.984	0.991	0.996	0.990	0.997	0.945	0.941
Time Chunks	12	12	12	12	12	12	12
Events	175,026	48,430	16,536	19,445	3,212	955	955
Belief Obs.	4,598,289	867,567	166,346	109,751	86,193	58,864	58,864
p-val: Const = 0	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
p-val: Slope = 1	<0.001	<0.001	<0.001	0.274	0.002	0.004	0.025

Notes: This table replicates Table III, but with 12 equal-length time windows, rather than 24. See that table's notes for details on estimation and interpretation.

Figure A7. Movement and Uncertainty Reduction for Sports Betting: 36 Time Chunks

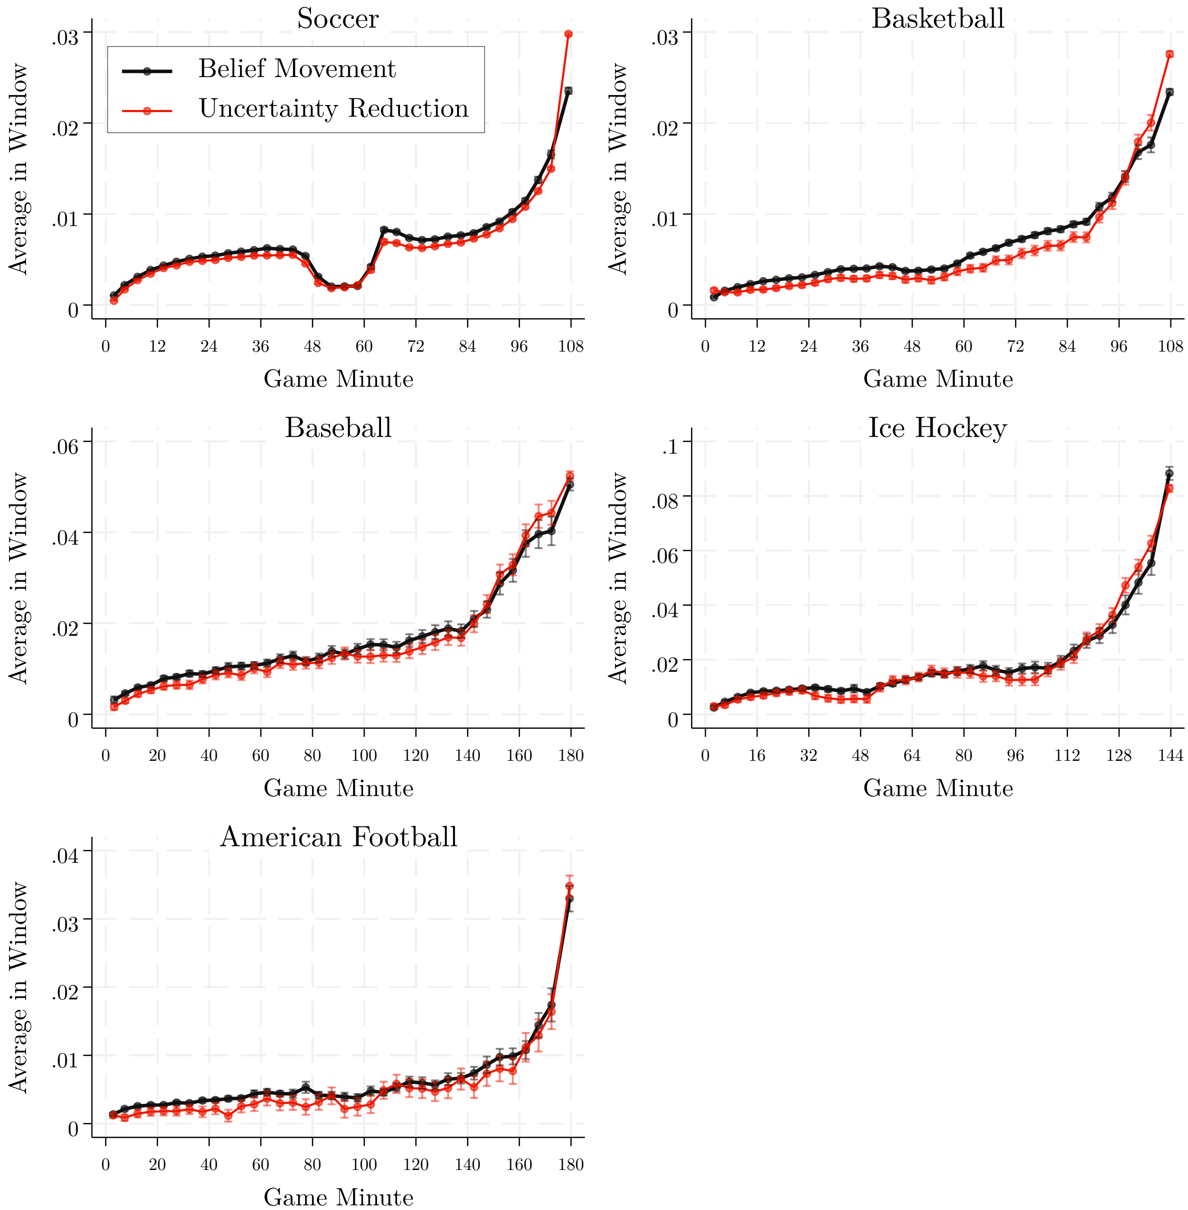

Notes: This figure replicates Figure VII, but with 36 equal-length time windows, rather than 24. See that figure's notes for details on construction.

Figure A8. Movement and Uncertainty Reduction for Finance Data: 36 Time Chunks

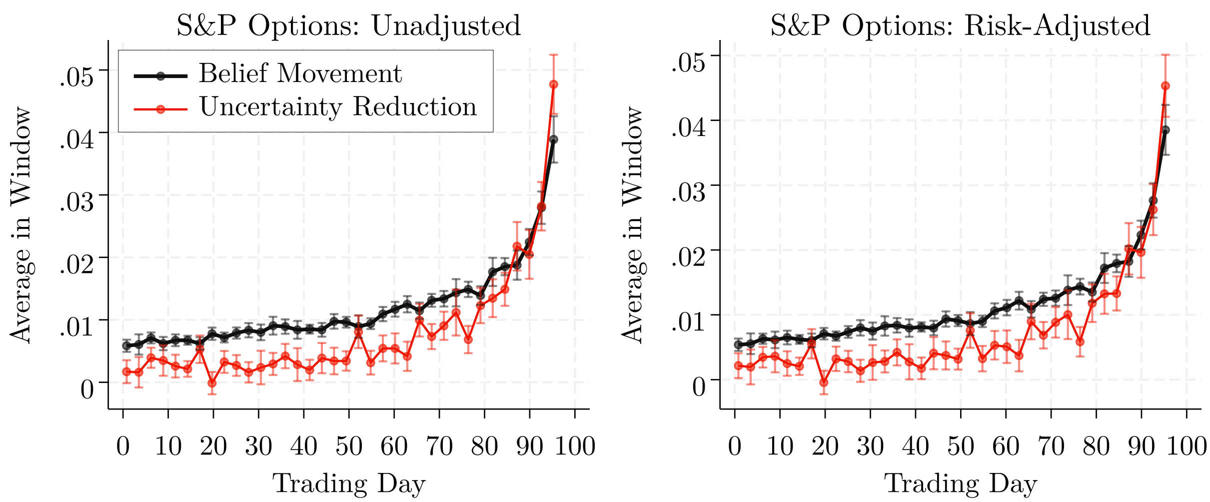

Notes: This figure replicates Figure VIII, but with 36 equal-length time windows, rather than 24. See that figure's notes for details on construction.

Table A4. Regressions of Movement on Uncertainty Reduction: 36 Time Chunks

Dep. Var.:	Sports					Finance	
Movement	Soccer	Basketball	Baseball	Hockey	Football	Raw	Risk-Adj.
Constant	0.0014	0.0016	0.0027	0.0020	0.0015	0.0063	0.0058
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0001)	(0.0003)	(0.0003)
Uncert. Red.	0.847	0.849	0.883	0.925	0.920	0.705	0.751
	(0.003)	(0.008)	(0.015)	(0.013)	(0.026)	(0.035)	(0.040)
R^2	0.955	0.974	0.993	0.975	0.982	0.932	0.928
Time Chunks	36	36	36	36	36	36	36
Events	175,026	48,430	16,536	19,445	3,212	955	955
Belief Obs.	4,598,289	867,567	166,346	109,751	86,193	58,864	58,864
p-val: Const = 0	<0.001	<0.001	<0.001	<0.001	0.051	<0.001	<0.001
p-val: Slope = 1	<0.001	<0.001	<0.001	<0.001	0.054	<0.001	<0.001

Notes: This table replicates Table III, but with 36 equal-length time windows, rather than 24. See that table's notes for details on estimation and interpretation.

Further Robustness Results for Options Data

FIGURE A9. Movement and Uncertainty Reduction for Options: Alternative Adjustments

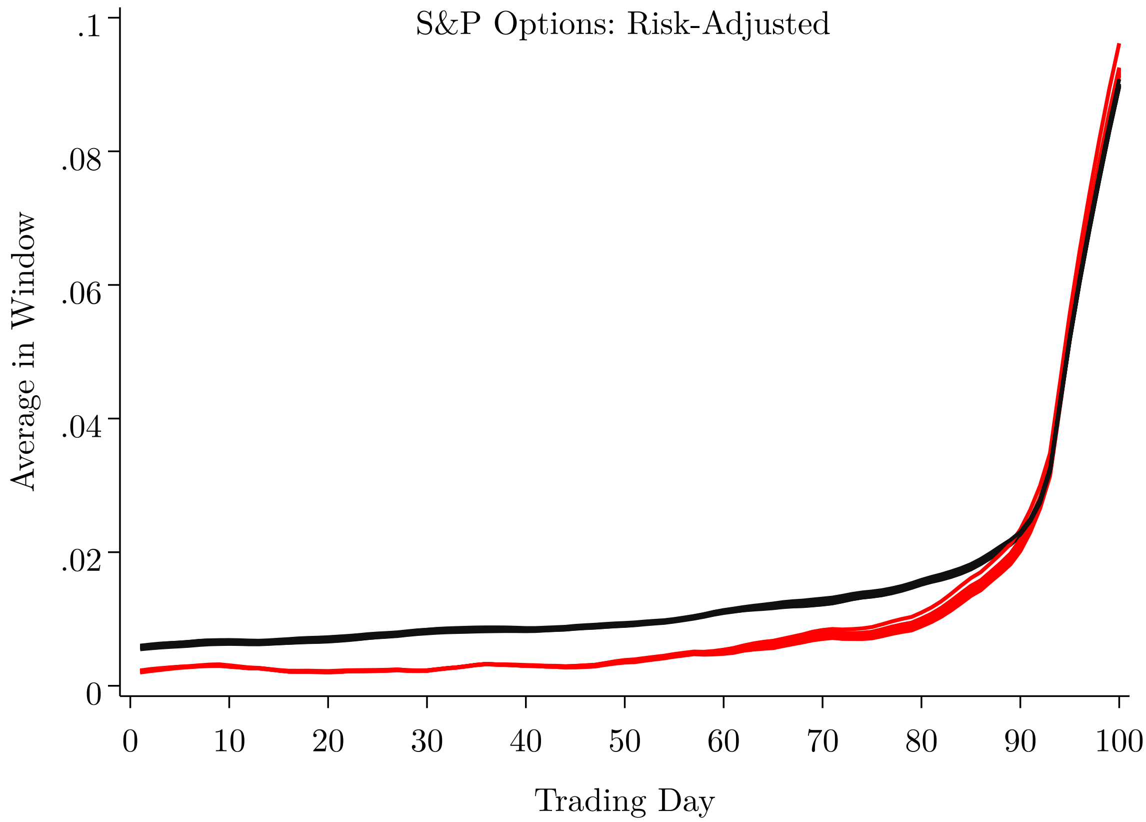

Notes: This figure replicates the right panel of Figure VIII to show the smoothed average movement (black lines) and uncertainty reduction (lighter red lines) statistics over time for the beliefs implied by option data, but with alternative risk adjustments. Each line represents a different method to calculate risk-adjusted beliefs from the raw, unadjusted risk-neutral beliefs, as described in Appendix B.2. Some aspects of the figure (including confidence intervals) are omitted to enable a clear view of the range of plotted lines across risk adjustments. While the different risk-adjustment methods do lead to different inferred beliefs, the broad pattern of movement and uncertainty curves is very similar across the methods, as the curves are close to overlapping in most cases.

FIGURE A10. Movement and Uncert. Red. for Finance: Binarized, Noise-Corrected Beliefs

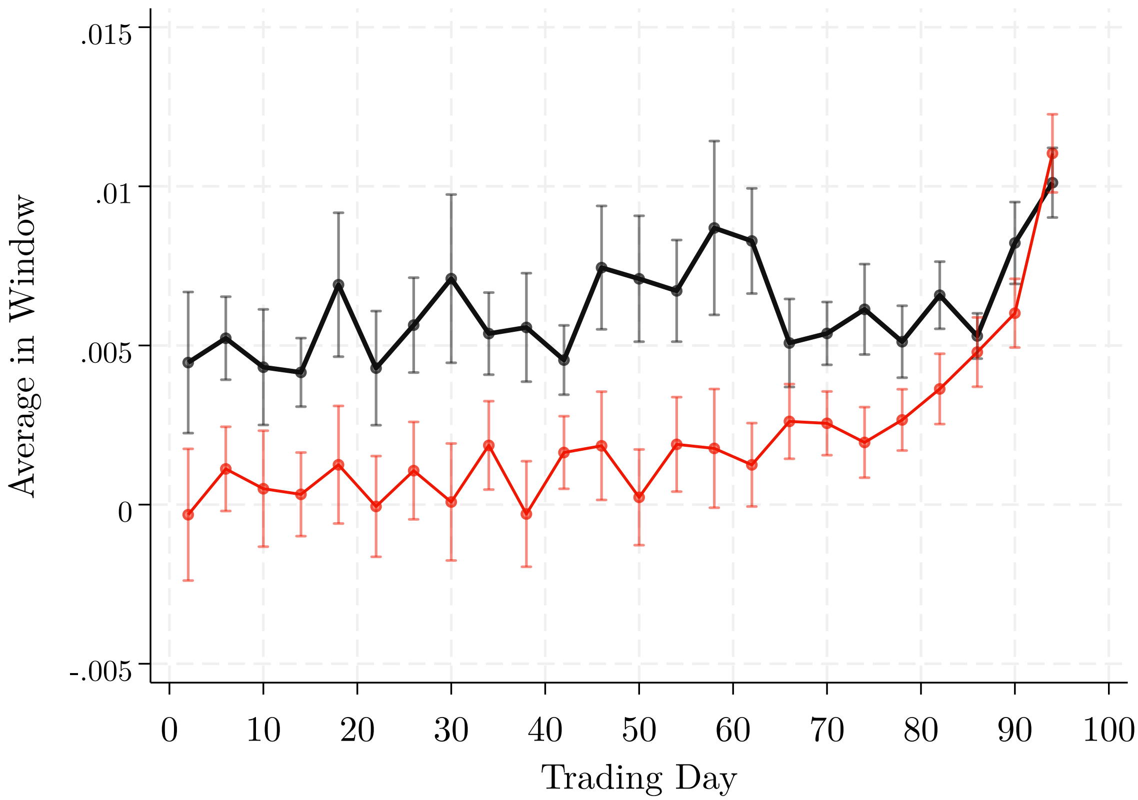

Notes: This figure replicates the left panel of Figure VIII, but using the binarized and noise-corrected risk-neutral beliefs data from Augenblick and Lazarus (2023). Belief movement is plotted in black, and uncertainty reduction in lighter red. Data are not adjusted for risk aversion. See Figure VIII for details on the plot, and see Section 6 and Online Appendix C.6 of AL (2023) for details on the noise correction and binarization.

FIGURE A11. Movement and Uncertainty Reduction for Finance Data: Post-2000

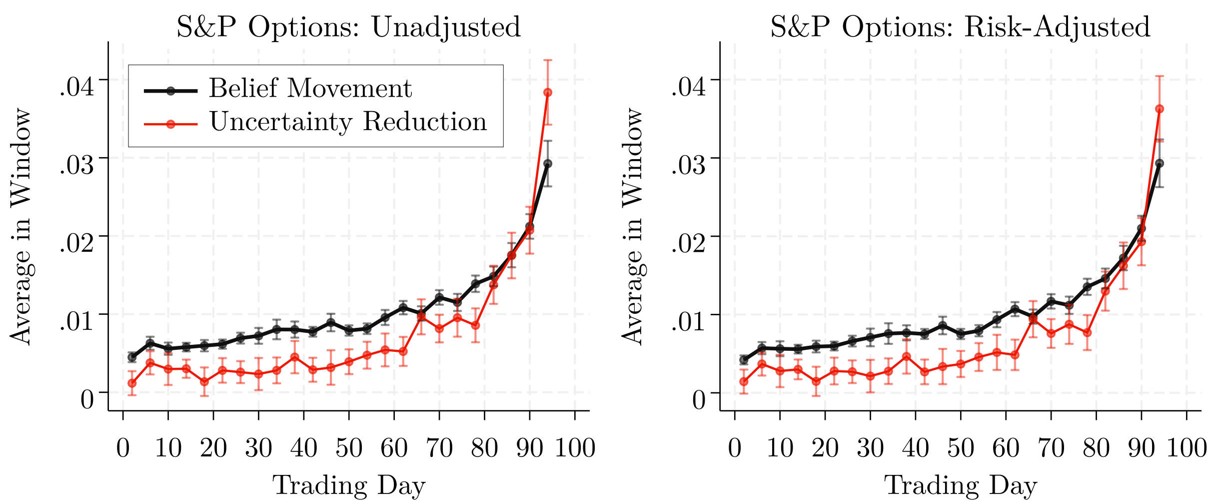

Notes: This figure replicates Figure VIII, using only data after the year 2000. See that figure's notes for details on construction. The figure demonstrates that our option results are robust to not including the early part of the sample, which contains somewhat noisier option data (see AL 2023).

FIGURE A12. Movement and Uncertainty Reduction for Finance Data: Post-2010

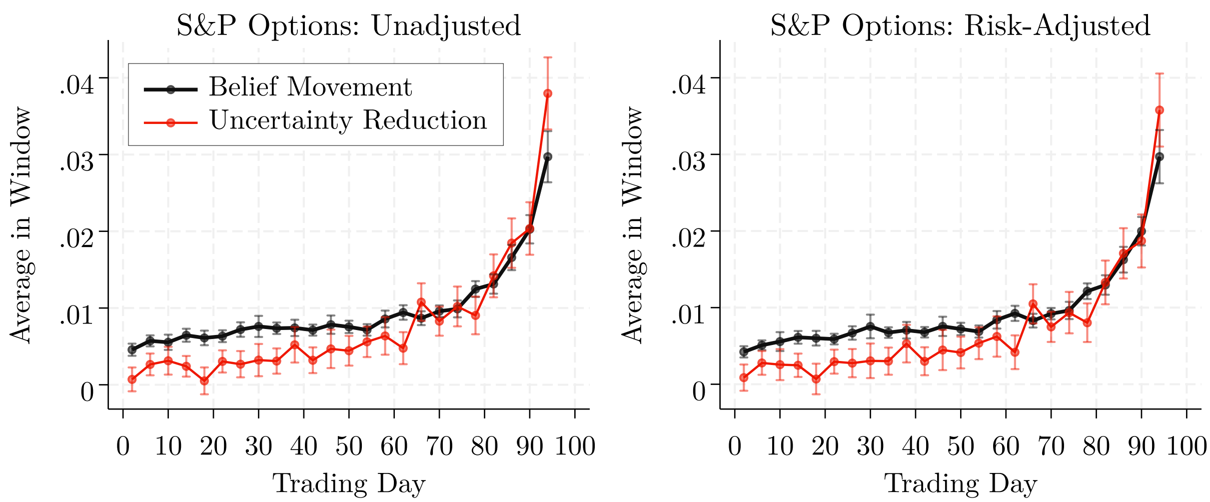

Notes: This figure replicates Figure VIII, using only data after 2010. See that figure's notes for details on construction. Along with Figure A11, this figure demonstrates that our option results are robust across subsamples.

C. Experiment Study Materials

C.1. Study 1a

Overview and Instructions

Princeton Experimental Laboratory for the Social Sciences

Overview and Bonus Payment

On the following pages, you will be asked to make a series of choices that can impact your bonus payment.

After all participants complete the study, 5 participants will be chosen at random to receive a bonus payment of up to $100 based on their choices. The high bonus is because it is important for us that you take this study seriously.

You will see between 25 and 35 question pages, most of which are similarly styled. These questions are divided into sections. There will be an "attention check" question in the study. The answer to this question will be obvious to anyone paying attention. Participants who do not answer the attention check question correctly will still receive their show-up payment, but **will not be eligible for a bonus payment**.

Instructions for the study are on the following page.

Instructions

On the following several pages, you will be asked to predict which jar a randomly-chosen ball comes from.

For instance, you will see questions like the following:

You draw a card from one of two modified decks of cards; a **Green** deck or a **Purple** deck.

The **Green** deck has **15 Diamonds** (♦) and **10 Spades** (♠).
The **Purple** deck has **10 Diamonds** (♦) and **15 Spades** (♠).

*The card you draw is a **Diamond (♦)**.*

What do you think is the percent chance that your **Diamond** (♦) came from the **Green** deck vs. the **Purple** deck?

Please answer between 0 and 100 for each question and have your answers sum up to 100, where higher numbers mean you think that deck is more likely. Your answer may include decimals.

Percent chance that your **Diamond** (♦) came from the **Green** deck:

_______ percent

Percent chance that your **Diamond** (♦) came from the **Purple** deck:

_______ percent

We have carefully chosen the payment rule so that you will earn the most bonus payment on average if you give a guess that you think is the true likelihood. If you are interested, further details on the payment rule are below.

Payment for your prediction:

If you are selected to receive a bonus payment, to determine your payment the computer will randomly choose a question and then randomly draw two numbers. For each draw, all numbers between 0 and 100 (including decimal numbers) are equally likely to be selected. Draws are independent in the sense that the outcome of the first draw in no way affects the outcome of the second draw.

If the selected Jar is Jar 1 and the number you picked is larger than either of the two draws, you will get a bonus payment of $100.

If the selected Jar is Jar 2 and the number you picked is smaller than either of the two draws, you will get a bonus payment of $100.

Otherwise, you will get a bonus payment of $10.

Main Decision Screen

Princeton Experimental Laboratory
for the Social Sciences

You are on Question 1 of 15 in this section.

You draw a card from one of two modified decks of cards; a **Green** deck or a **Purple** deck.

The **Green** deck has **853 Diamonds** (♦) and **812 Spades** (♠).
The **Purple** deck has **812 Diamonds** (♦) and **853 Spades** (♠).

*The card you draw is a **Spade** (♠).*

What do you think is the percent chance that your **Spade** (♠) came from the **Green** deck vs. the **Purple** deck?

Please answer between 0 and 100 for each question and have your answers sum up to 100, where higher numbers mean you think that deck is more likely. Your answer may include decimals.

Percent chance that your **Spade** (♠) came from the **Green** deck	0	percent
Percent chance that your **Spade** (♠) came from the **Purple** deck	0	percent
Total	0	percent

Attention Check

PEXL
Princeton Experimental Laboratory for the Social Sciences

You draw a card from one of two modified decks of cards; a **Green** deck or a **Purple** deck.

The **Green** deck has **867 Diamonds** (♦) and **5309 Spades** (♠).
The **Purple** deck has **5309 Diamonds** (♦) and **867 Spades** (♠).

*The card you draw is a **Spade** (♠).*

This question is not like the previous ones. It is a check to make sure you are paying attention. Instead of answering the question normally, please answer 3 for the first question and 14 for the second one. You will not be eligible for a bonus payment if you get this question incorrect.

Percent chance that your **Spade** (♠) came from the **Green** deck	0 percent
Percent chance that your **Spade** (♠) came from the **Purple** deck	0 percent
Total	0 percent

Unknown Signal Strength

Princeton Experimental Laboratory for the Social Sciences

You draw a card from one of two modified decks of cards; a **Green** deck or a **Purple** deck.

This time, you do not know the composition of the two decks for sure.

The **Green** deck **either** has:

- **841 Diamonds** (♦) and **824 Spades** (♠)

  OR

- **448 Diamonds** (♦) and **1217 Spades** (♠).

The **Purple** deck **either** has:

- **824 Diamonds** (♦) and **841 Spades** (♠)

  OR

- **1217 Diamonds** (♦) and **448 Spades** (♠).

Both compositions are equally likely.

*The card you draw is a **Diamond** (♦).*

What do you think is the percent chance that your **Diamond** (♦) came from the **Green** deck vs. the **Purple** deck?

Please answer between 0 and 100 for each question and have your answers sum up to 100, where higher numbers mean you think that deck is more likely. Your answer may include decimals.

Percent chance that your **Diamond** (♦) came from the **Green** deck	0 percent
Percent chance that your **Diamond** (♦) came from the **Purple** deck	0 percent
Total	0 percent

Princeton Experimental Laboratory
for the Social Sciences

You are on Question 1 of 5 in this section.

On the previous questions, you were given a card and asked to predict what deck it came from.

On this question, you will choose *how many* cards to draw in order to help you with your predictions. Every card is drawn with replacement.

You may draw up to 3 cards at a cost. If you are selected for the bonus and this round is chosen for payment, you can win up to $100. The first card costs $1, the second card costs an additional $2 ($3 total), and the third card costs an additional $3 ($6 total).

- $1 roughly corresponds to moving from being 50% sure that you know the deck to being 60% sure;
- $3 corresponds to moving from being 50% sure that you know the deck to being 67% sure;
- $6 corresponds to moving from being 50% sure that you know the deck to being 74% sure.

If you think the cards are helpful in distinguishing the **Green** and **Purple** decks, you should draw more cards. If you think the cards are unhelpful in distinguishing the **Green** and **Purple** decks, you should not draw cards.

The **Green** deck has **853 Diamonds** (♦) and **812 Spades** (♠).
The **Purple** deck has **812 Diamonds** (♦) and **853 Spades** (♠).

You can now choose how many cards to draw from the deck.

- ◯ Do not draw any cards
- ◯ Draw 1 card (Total cost: 1 point)
- ◯ Draw 2 cards (Total cost: 3 points)
- ◯ Draw 3 cards (Total cost: 6 points)

The **Green** deck has **853 Diamonds** (♦) and **812 Spades** (♠).
The **Purple** deck has **812 Diamonds** (♦) and **853 Spades** (♠).

You chose to draw 2 cards.

What do you think the percent chance (between 0 and 100) is that the card is from the **Green** deck if the cards drawn are:

0	2 Spades (♠)
0	2 Diamonds (♦)
0	1 Diamond (♦) and 1 Spade (♠)

→

Demographics

Princeton Experimental Laboratory for the Social Sciences

What is your age?

What is your gender?

- Female
- Male
- Other / Prefer not to answer

In politics today, do you consider yourself a Republican, a Democrat, or an Independent?

- Democrat
- Republican
- Independent

What is your highest level of education?

- Did not graduate high school
- High school graduate, diploma, or equivalent (such as GED)
- Began college, no degree
- Associate's degree
- Bachelor's degree
- Postgraduate or professional degree

What is your race/ethnicity?

- ◯ American Indian
- ◯ White
- ◯ Asian
- ◯ Black or African American
- ◯ Hispanic or Latino
- ◯ Two or more of these
- ◯ Other / Prefer not to answer

Cognitive Reflection Test

PEXL
Princeton Experimental Laboratory
for the Social Sciences

Quiz 1

A bat and a ball cost $10.50 in total. The bat costs $10.00 more than the ball.

How much does the ball cost?

- ◯ $0.50
- ◯ $10.00
- ◯ $10.25
- ◯ $0.25

Quiz 2

If it takes 7 machines 7 minutes to make 7 widgets, how long would it take 70 machines to make 70 widgets?

- ◯ 0.1 minutes
- ◯ 490 minutes
- ◯ 70 minutes
- ◯ 7 minutes

Quiz 3

In a community, there is a rapidly-spreading virus. Every day, the virus infects twice as many people. If it takes 42 days for the virus to infect the entire community, how long would it take the virus to infect half the community?

- ◯ 21 days
- ◯ 42 days
- ◯ 2 days
- ◯ 41 days

C.2. Study 1b

Overview and Instructions

Overview and Bonus Payment

On the following pages, you will be asked to make a series of choices that can impact your bonus payment.

After all participants complete the study, 10 participants will be chosen at random to receive a bonus payment of up to $50 based on their choices. The high bonus is because it is important for us that you take this study seriously.

You will see 15 question pages, most of which are similarly styled. There will be an "attention check" question in the study. The answer to this question will be obvious to anyone paying attention. Participants who do not answer the attention check question correctly will still receive their show-up payment, but **will not be eligible for a bonus payment**.

Instructions for the study are on the following page.

Instructions

On the following several pages, you will be asked to predict which deck a card comes from. We will choose both the deck and the card at random. We have carefully chosen the payment rule so that you will earn the most bonus payment on average if you give a guess that you think is the true likelihood. If you are interested, further details on the payment rule are below.

You will see questions that look like the following:

*We will randomly select one of three modified decks of cards --- either the **Green** deck, the **Light blue** deck, or the **Dark blue** deck --- and you will draw one card from one of the decks:*

*The **Green** deck has **15 Diamonds** and **10 Spades**.*
*The **Light blue** deck has **10 Diamonds** and **15 Spades**.*
*The **Dark blue** deck has **10 Diamonds** and **15 Spades**.*

*The card you draw is a **Diamond**.*

*What do you think is the percent chance that your **Diamond** came from the **Green** deck, or each of the **Blue** decks?*

*Please answer between 0 and 100 for each question and have your answers sum up to 100, where higher numbers mean you think that deck is more likely. Your answer may include decimals. Since both **Blue** decks are identical, your answer should be the same for each of these.*

*Percent chance that your **Diamond** came from the **Green** deck: ______%*
*Percent chance that your **Diamond** came from the **Light blue** deck: ______%*
*Percent chance that your **Diamond** came from the **Dark blue** deck: ______%*

Payment rule:

If you are selected to receive a bonus payment, to determine your payment the computer will randomly choose a question and then randomly draw two numbers. For each draw, all numbers between 0 and 100 (including decimal numbers) are equally likely to be selected. Draws are independent in the sense that the outcome of the first draw in no way affects the outcome of the second draw.

If the selected Deck is the Green Deck and the number you picked for the Green Deck is larger than either of the two draws, you will get a bonus payment of $50.

If the selected Deck is one of the Blue Decks and the sum of the numbers you picked on Blue Decks is smaller than either of the two draws, you will get a bonus payment of $50.

Otherwise, you will get a bonus payment of $0.

Three Decks

For the next five questions, you will see cards drawn from one of **three decks**:

- The **Green** deck,
- The **Light blue** deck, or
- The **Dark blue** deck.

The two **Blue** decks will have the same numbers of Spades and Diamonds, while the **Green** deck will look different.

Main Decision Screen

Prediction

You are on Prediction 5 of 15.

We will randomly select one of three modified decks of cards --- either the **Green** deck, the **Light blue** deck, or the **Dark blue** deck --- and you will draw one card from one of the decks:

The **Green** deck has **629 Diamonds** and **1036 Spades**.
The **Light blue** deck has **1036 Diamonds** and **629 Spades**.
The **Dark blue** deck has **1036 Diamonds** and **629 Spades**.

The card you draw is a **Diamond**.

What do you think is the percent chance that your **Diamond** came from the **Green** deck, or each of the **Blue** decks?

Please answer between 0 and 100 for each question and have your answers sum up to 100, where higher numbers mean you think that deck is more likely. Your answer may include decimals. Since both **Blue** decks are identical, your answer should be the same for each of these.

Percent chance that your **Diamond** came from the **Green** deck	0	%
Percent chance that your **Diamond** came from the **Light blue** deck	0	%
Percent chance that your **Diamond** came from the **Dark blue** deck	0	%
Total	0	%

Attention Check

This question is not like the previous ones; it is a check to make sure you are paying attention. You will not be eligible for a bonus payment if you get this question incorrect.

You draw a card from one of two modified decks of cards; a **Green** deck or a **Blue** deck.

The **Green** deck has **867 Diamonds (♦)** and **5309 Spades (♠)**.
The **Blue** deck has **5309 Diamonds (♦)** and **867 Spades (♠)**.

Instead of answering the question normally, please answer 3 for the first question and 14 for the second one.

What is your answer to the **first** question?	0	%
What is your answer to the **second** question?	0	%
Total	0	%

Previous guess

You are on Prediction 5 of 15.

Your decision on the previous screen indicates that you believe there is a **22% chance** that your **Diamond** came from the **Green** deck.

How certain are you that the optimal guess is somewhere between **21%** and **23%**?

Very uncertain **Completely certain**

0 10 20 30 40 50 60 70 80 90 100

C.3. Study 2

Overview and Instructions

Overview and Bonus Payment

On the following pages, you will be asked to make a series of choices that can impact your bonus payment.

After all participants complete the study, 10 participants will be chosen at random to receive a bonus payment of up to $50 based on their choices. The high bonus is because it is important for us that you take this study seriously.

You will see 16 question pages, most of which are similarly styled. There also will be an "attention check" question in the study. The question will look similar to other questions, but the answer will be obvious to anyone paying attention. Participants who do not answer the attention check question correctly will still receive their show-up payment, but **will not be eligible for a bonus payment**.

Instructions for the study are on the following page.

Instructions

Over the course of this study, we will give you information about various scenarios in NBA games, and ask you to predict the likelihood that each team wins. We will not tell you the names of the teams. We have used a model based on a database of regular-season NBA games with several years of play-by-play data to estimate the likelihoods of each team winning in these scenarios.

The closer your answer is to the likelihood, the more likely you are to win the $50 bonus. The exact way that your bonus payment is determined may seem complicated, but what is important to know is that these rules were carefully determined so that you have the best chance of winning the $50 bonus if you honestly and carefully answer these questions. If you are interested, further details on the payment rule are at the bottom.

Below is an example of a situation you might see:

Prediction

We will give you some details of an NBA game and ask you to predict the likelihood that Team X or Team Y wins the game.

With 6:15 left in the second quarter, Team X made a shot to take the lead by 4 points.

Time	**Last action**	**Score**	**Win probability**
Q2: 6:15 left	**Team X made a shot**	**Team X is up by 4 points**	**?**

What do you think is the percent chance that Team X wins the game?

Payment for your prediction:

If you are selected to receive a bonus payment, to determine your payment the computer will randomly choose a question and then randomly draw two numbers. For each draw, all numbers between 0 and 100 (including decimal numbers) are equally likely to be selected. Draws are independent in the sense that the outcome of the first draw in no way affects the outcome of the second draw.

If ***Team X won*** *and the number you picked (as the chance of the team winning) is* ***larger*** *than either of the two draws, you will get a bonus payment of $50.*

If ***Team Y won*** *and the number you picked (as the chance of the team winning) is* ***smaller*** *than either of the two draws, you will get a bonus payment of $50.*

Otherwise, you will not receive a bonus payment.

Main Decision Screen

Prediction

We will give you some more information about the game. With 2:25 left in the third quarter, Team F missed a shot and Team E rebounded the ball.

Time	Last action	Score	Your guess
Q3: 2:40 left	Team E made a shot	Team E is up by 5 points	Team E: **66%** chance of winning
Q3: 2:25 left	**Team F missed a shot, Team E rebounded the ball**	**Team E is up by 5 points**	?

Now, what do you think is the percent chance that Team E wins the game?

0 10 20 30 40 50 60 70 80 90 100

Please answer between 0 and 100, where higher numbers mean you think Team E is more likely to win. Your payment will be determined based on how close your answer is to the actual estimated likelihood that Team E wins.

Attention Check

Prediction

We will give you some more information about the game, but this question is not like the previous ones; it is a check to make sure you are paying attention. With 1:50 left in the second quarter, Team C missed a shot and Team D rebounded the ball.

Time	**Last action**	**Score**	**Your guess**
Q2: 2:40 left	Team C made a shot	Team C is up by 5 points	Team C: **61%** chance of winning
Q2: 2:25 left	Team D made a shot	Team C is up by 3 points	Team C: **58%** chance of winning
Q2: 2:10 left	Team C made a shot	Team C is up by 5 points	Team C: **62%** chance of winning
Q2: 2:00 left	Team D made a shot	Team C is up by 3 points	Team C: **57%** chance of winning
Q2: 1:50 left	**Team C missed a shot, Team D rebounded the ball**	**Team C is up by 3 points**	?

What quarter is the game in?

You will not be eligible for a bonus payment if you get this question incorrect.