Gain control of saccadic eye movements is probabilistic

Matteo Lisi[a], **Joshua A. Solomon**[b], **and Michael J. Morgan**[b]

[a]Department of Biological and Experimental Psychology, Queen Mary University of London, E1 4NS, UK; [b]Centre for Applied Vision Research, City, University of London, London EC1V 0HB, UK

This manuscript was compiled on June 14, 2019

Saccades are rapid eye movements that orient the visual axis toward objects of interest to allow their processing by the central, high-acuity retina. Our ability to collect visual information efficiently relies on saccadic accuracy, which is limited by a combination of uncertainty in the location of the target and motor noise. It has been observed that saccades have a systematic tendency to fall short of their intended targets, and it has been suggested that this bias originates from a cost function that overly penalizes hypermetric errors. Here we tested this hypothesis by systematically manipulating the positional uncertainty of saccadic targets. We found that increasing uncertainty produced not only a larger spread of the saccadic endpoints but also more hypometric errors and a systematic bias toward the average of target locations in a given block, revealing that prior knowledge was integrated into saccadic planning. Moreover, by examining how variability and bias co-varied across conditions, we estimated the asymmetry of the cost function and found that it was related to individual differences in the additional time needed to program secondary saccades for correcting hypermetric errors, relative to hypometric ones. Taken together, these findings reveal that the saccadic system uses a probabilistic-Bayesian control strategy to compensate for uncertainty in a statistically principled way and to minimize the expected cost of saccadic errors.

Motor control | Cost function | Eye movements | Saccades

Saccadic eye movements serve a pivotal role in the foveate visual systems of primates, by quickly orienting the fovea (the central, high-acuity part of the retina) toward objects of interest. It seems reasonable to surmise that saccades have evolved to serve vision optimally, however it is not obvious what the optimum should be. Given that visual sensitivity is much reduced during saccades, one relevant cost to minimize could be the time spent in-flight. However, as it has been pointed out (1), duration cannot be the only factor, otherwise oblique saccades should be significantly faster than purely horizontal or vertical ones, and they are not (2). Another crucial factor is accuracy: like all our movements, saccades are variable and often miss the desired destination due to motor noise and sensory uncertainty. These errors might have undesirable consequences, such as hindering the timely identification of dangers in the environment. Indeed, it has been shown that the stereotypical kinematics of saccadic eye movements (the so-called 'main sequence') are optimal for minimizing the variability (and thus the mean error) of landing positions in the presence of signal-dependent motor noise (3). In light of this, it may seem surprising that, on top of their inescapable variability, saccades display a systematic hypometric bias: they tend to fall short of their target by a fixed proportion of the target distance, about 10% (4).

What is the origin of this bias, and why has evolution not corrected it? One possible explanation relates to the programming of secondary saccades, which are often needed to correct the saccadic landing error. Importantly, the time required to launch these corrective saccades is not independent of the error of the initial primary saccade: corrective saccades are slower to launch when they are in the opposite direction relative to the primary saccade (5–7). If the total time needed to reach the desired destination (including the latency of corrections) were a relevant factor, then the ideal strategy would be to plan saccades that are, on average, hypometric, thereby decreasing the relative likelihood of overshoot errors. Formally, this can be expressed with an asymmetrical cost function, i.e. one that assigns a greater cost to an overshoot error relative to an undershoot of the same magnitude. Although this strategic account of saccadic hypometria is appealing, it lacks direct empirical support. In fact, other studies have proposed the alternative view that undershoots may be best viewed as an inevitable property of the oculomotor system (8), due to sub-optimal sensorimotor transformations.

Assuming that biases in saccadic targeting are due to a deliberate strategy and this strategy is probabilistic (i.e. if it accounts for uncertainty in a statistically principled way) and Bayesian, two predictions can be made. First, variability and bias should be systematically related one another and the ideal saccadic gain (the ratio of saccadic amplitude and target distance) should decrease when uncertainty about the position of the target increases, as demonstrated in Fig. 1 (see figure caption for detailed explanation). Second, if the strategy is Bayesian it should take advantage of prior information whenever available. Results consistent with this latter prediction have been reported by Kapoula and Robinson (9, 10),

Significance Statement

Decades of research have shown that, when measured in laboratory conditions, saccadic eye movements are not only variable, due to noise in sensory-motor pathways, but also inaccurate, displaying systematic biases toward smaller movement amplitudes (hypometria) or toward the mean location of the targets (central tendency). Here, we show that such biases are explained by a probabilistic strategy to find the optimal balance of bias and variance that minimizes the expected costs of saccadic errors. Our findings indicate that the oculomotor system possesses a probabilistic representation of its own sensorimotor uncertainty and uses that representation to adjust the parameters of each saccade.

M.L., M.J.M., and J.A.S. designed research; M.L. performed research; M.L. analyzed data; M.L. wrote the paper with input from all authors.

The authors declare no competing financial interests.

To whom correspondence should be addressed. E-mail: matteo.lisi@inventati.org

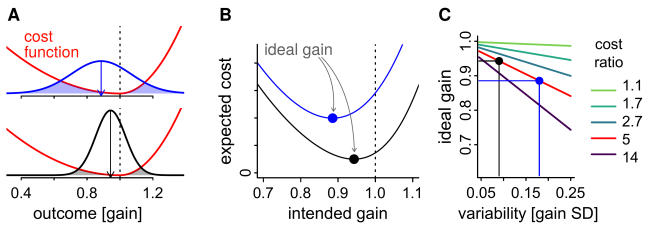

Fig. 1. Predicted relationship between saccadic variability and undershoot. A. The red curve represents the cost of a saccadic error plotted against gain (proportion of target distance). The two Gaussian curves represent the expected distributions of motor outcomes for two conditions with different uncertainties about the location of the target: in the condition with larger uncertainty (blue) there is a broader range of motor outcomes for a given motor command (intended gain, represented by the vertical arrow). The expected cost for a certain intended gain is computed by integrating all possible motor outcomes, weighted by their probabilities. B. The expected cost is plotted as a function of the intended gain. When uncertainty is larger the expected cost is overall higher, and the ideal gain (which minimizes the expected cost) shifts toward more hypometric values. C. Relationship between ideal gain and saccadic endpoint variability, for different degrees of asymmetry. The asymmetry is quantified as the ratio between the cost of an overshoot relative to an undershoot of the same size. Since the asymmetry determines the slope of the relationship between gain and variability, it is possible to estimate it by measuring (at least) two different conditions with varying levels of uncertainty.

who found that saccades display a range effect, i.e. a bias toward the mean of target positions in a given block. Results seemingly inconsistent with this latter prediction appear in two recent studies, which failed to replicate the range effect (11, 12); however these studies did not manipulate uncertainty. Here we aimed to assess whether a range effect would appear as uncertainty increased. Indeed, any central tendency bias (13) arising from a probabilistic combination of sensory likelihood and prior knowledge should increase as the likelihood becomes more diffuse.

Results

In order to test the two predictions mentioned in the Introduction, we conducted a series of experiments in which we manipulated the positional uncertainty of the saccadic target, as well as the range of its possible positions (thus their prior probabilities), and measured how these factors contribute to constant and variable saccadic errors. We were interested in simple visual orienting responses, therefore we avoided adding more explicit tasks that may have influenced the cost function. We expected both the hypometric bias and the range effect to increase with increasing uncertainty. In Experiment 1 (n=12) we manipulated the uncertainty by blurring a Gaussian blob embedded in noise (keeping the total luminance energy constant, see Fig. 2A), and measured saccadic responses in two sessions, run on separate days, that contained different ranges of target eccentricities (this was necessary to measure the range effect). Although positional uncertainty should be reflected in the distribution of saccade endpoints, to make sure that our manipulation was successful, we also measured each observer's perceptual precision for comparing the eccentricities of blurred targets in a purely psychophysical task. The results confirmed that blurring the targets increases the uncertainty of judgments about their positions (see SI). To characterize further the relationship between sensory uncertainty and saccadic targeting, we conducted two additional experiments. In Experiment 2 (n=20) we varied independently the size and the peak luminance of the saccadic target (Fig. 2A). This experiment determined the relative contributions of pure changes in target size and visibility. In Experiment 3 (n=26) we further investigated the robustness of the saccadic range effect, by running the two sessions in the same day and using targets that varied only in visibility (but not size). Since these experiments provide complementary findings, in the following we report the results organized by thematic points. Detailed information about experimental procedures and statistical modelling is reported in the SI.

Positional uncertainty increases saccadic variability and hypometria. We found that increasing the space constant of a Gaussian blob increased the variability of the amplitudes of saccades directed to it, $F(2, 22) = 5.66, p = 0.01$. Crucially, we found that greater uncertainty not only increased the variable error, but also the undershoot (see Fig. 2B). We assessed the variations of saccadic undershoot by means of a multilevel (mixed-effects) linear model (see SI for details), with saccadic amplitude as dependent variable and target distance and blob's σ as predictors. The estimates of model parameters indicate that the saccadic gain (the slope of the linear relationship between saccadic amplitudes and target distance) was already hypometric in the condition with smallest σ (baseline gain 0.93±0.06, mean ± standard error) and became even more hypometric as σ increased: the differences from baseline were -0.01±0.03, for the condition with σ=0.9deg; and -0.17±0.03, for the condition with σ=1.5deg. The finding of a simultaneous increase in variable and constant errors is to be expected under the hypothesis of an asymmetrical cost function (Fig. 1). Moreover, the total changes in variability and bias (quantified as the difference between the condition with largest and smallest uncertainty) were correlated across participants (Pearson's r=-0.73, 95%CI [-0.92, -0.23]): participants who showed the largest increase in endpoint variability also displayed the largest decrease in saccadic gain, suggesting a systematic relationship between variability and bias.

The blur manipulation used in Experiment 1 simultaneously decreased the target's peak luminance, and increased its size. Saccades might have been biased toward the nearest edge of the target (e.g. the nearest zero-crossing in the second derivative or perhaps the half-height of the luminance profile (14)). The relative contributions of visibility and size could not have been distinguished within Experiment 1, so we designed Experiment 2 to discriminate between them. The procedure was similar, however we varied the stimuli in two distinct conditions. In the first condition size (σ) was kept constant, while we varied the peak luminance (fixed-size; Fig. 2A); this condition was designed to measure how visibility and signal-to-noise ratio affect saccadic eye movements when size is kept constant. In the second condition we kept luminance fixed at its maximum value, removed the background noise (minimizing the possible sources of uncertainty), and varied the size (σ) of the blobs (fixed-luminance); this condition was designed to isolate modulations of saccadic movements that were due only to variations of target size.

We found that both manipulations increased the variability of saccadic gain: fixed-luminance, $F(2, 38) = 11.29$, $p = 1.42 \times 10^{-4}$; fixed-size, $F(2, 38) = 16.84$, $p = 5.8 \times 10^{-6}$. Variability however increased up to higher levels in the fixed-size than in the fixed-luminance condition, $t(19) = 3.51$, $p = 0.002$. In both conditions, the increase in variability was accompanied by a decrease in saccadic amplitudes, albeit with some qualitatively

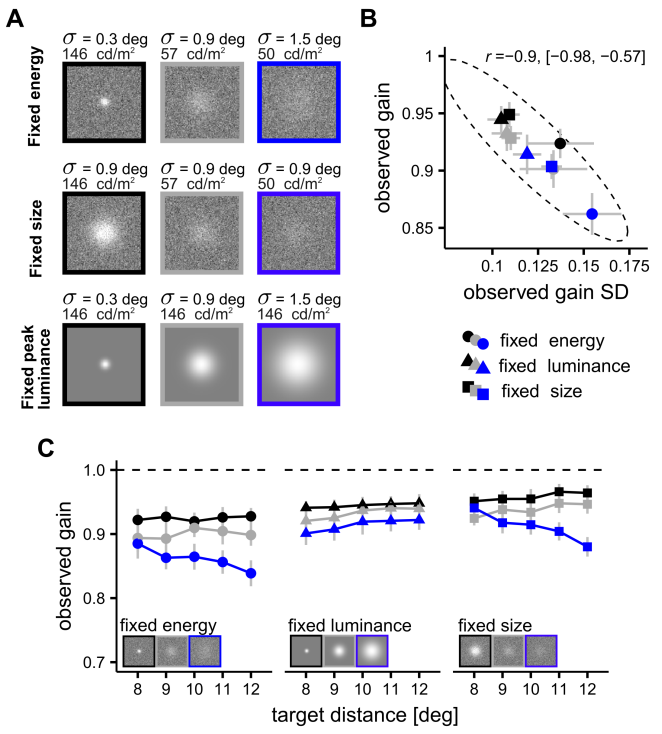

Fig. 2. Manipulation of positional uncertainty increases both behavioral variability and saccadic undershoot. A. Example of the stimuli used (see main Text and SI for details). B. Empirical relationship between variability and gain; each symbol represents the weighted average values (i.e. across observers) for the mean and standard deviation of saccadic gain in one experimental condition. Saccadic gain is negatively correlated with saccadic variability, as predicted by the theory (Fig. 1). C. Saccadic gain, plotted as a function of target distance (Experiment 1 and 2), for three different manipulations of the saccadic target. Only when the luminance is varied (fixed energy and fixed size conditions) does the decrease in amplitude vary as a function of target distance, suggesting the presence of a central bias. All error bars are bootstrapped standard errors.

different features. To quantify these features, we fit the data from each condition with a multilevel (mixed-effects) linear model, which had saccadic amplitude as a dependent variable and target distance and uncertainty level (indexed either by the blob's σ or its peak luminance) as predictors. In the fixed-luminance condition, the decrease in amplitude was constant with respect to the distance of the target, so that the slope of the linear relationship between saccadic amplitude and target distance did not vary systematically with the value of σ, $\chi^2(2) = 0.66$, $p = 0.72$. Analysis of the fixed-size condition instead revealed a different pattern. We found that, relative to the baseline where the peak luminance was 146cd/m^2, the decrease in saccadic amplitude was not uniform across target distances, as indicated by a significant interaction between distance and luminance, $\chi^2(2) = 30.06$, $p = 2.96 \times 10^{-7}$. This result indicates that the decrease in saccadic gain was modulated by the eccentricity of the target: gain decreased more when eccentricity was larger (see Fig. 2C). This finding suggests a bias toward intermediate eccentricities contingent on the visibility of the target, corresponding to the range-effect mentioned in the Introduction (9, 10) (see next section).

Saccadic range-effect depends on positional uncertainty. In Experiments 1 and 3, each participant was tested under two different conditions, with different ranges of target eccentricity (Fig. 3). Here we analyzed the effect of the eccentricity range ('large' vs 'small' eccentricity range) on saccadic behavior. We started by examining how saccades made toward the intermediate targets (present in both ranges) were influenced by the session. In agreement with recent reports (11, 12), we found no evidence for a central tendency bias when uncertainty was smallest (σ=0.3 or luminance 146cd/m^2), as indicated by the absence of systematic differences between saccadic amplitudes directed toward the intermediate targets [Exp. 1, $t(11) = 0.59$, $p = 0.57$; Exp. 3, $t(11) = 0.37$, $p = 0.71$]. However, analogous differences varied systematically across conditions with different uncertainties, as indicated by a significant interaction between range and uncertainty level: Exp. 1, $F(1, 23) = 15.05$, $p = 7.59 \times 10^{-4}$; Exp. 3, $F(1, 23) = 15.05$, $p = 0.01$ (two-way repeated measures ANOVA).

In order to quantify more precisely the range effect using all saccades (and not only those directed at the intermediate target) we assumed that the effect was due to a compression of saccadic responses toward the mean of target eccentricity in the block (a form for central tendency bias) and estimated the amount of compression using a linear regression approach. The regression model can be expressed as $\hat{S}_i = \beta_0 + \beta_1[\alpha\bar{E} + (1-\alpha)E_i]$, where $\hat{S}_i$ and E_i are the predicted saccadic amplitude and the target eccentricity at trial i, $\bar{E}$ is the average eccentricity in the current session, and α is a weighting parameter. Positive values of α indicate a bias toward the mean eccentricity, quantified as the proportion of compression, such that a value of $\alpha = 1$ would indicate that all saccades targeted the same central location, regardless of the trial-by-trial target eccentricities. All parameters were allowed to vary across conditions with different σ. We estimated a Bayesian mixed-effects version of this model, with participant as grouping factor (see SI for details). We calculated 95% credible intervals for the fixed-effect estimates of the weighting parameter α and found that the amount of compression differed significantly from zero only in the condition with largest uncertainty: Experiment 1, $\sigma = 1.5$, $\alpha = 0.18$, 95%CI [0.06, 0.30]; Experiment 3, peak luminance 50cd/m^2, $\alpha = 0.09$, 95%CI [0.01, 0.17] (Fig. 3B). Thus, our results indicate that although a range effect is not normally present for small, highly visible targets, a systematic bias toward the mean eccentricity nonetheless emerges when uncertainty increases.

Cost asymmetry determines the relationship between saccadic variability and bias. We suggest that the observed modulations of saccadic gain are a consequence of the oculomotor system seeking to minimize a cost function, in which overshoots and undershoots are given different weights. If an asymmetrical cost function were underlying the relationship between saccadic variability and undershoot, then it should be possible to estimate the degree of asymmetry, as shown in Fig. 1. In order to simplify the analysis, we transformed saccadic amplitudes in gain values (proportions of target distance) and pooled data from different target eccentricities together. This allowed us to specify a unique cost function for all eccentricities, where the error is defined in gain units. We assumed that cost would be well approximated by a quadratic function of the error, augmented with an additional asymmetry term that set a fixed ratio between the cost of undershoot and overshoot errors (see SI for details). Maximum likelihood estimates of the asymmetry parameter indicate that participants behaved as if they were optimizing an asymmetrical cost function where

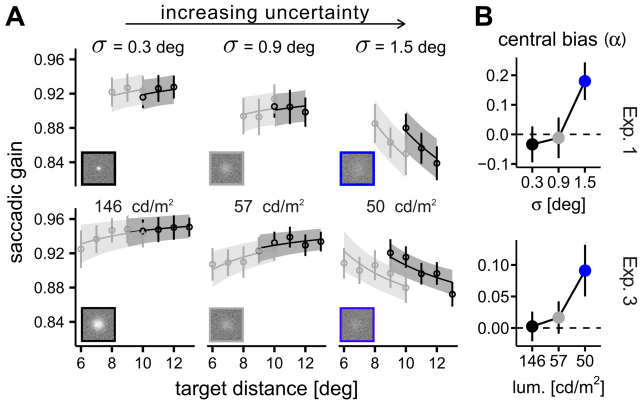

Fig. 3. The range effect. A. Mean saccadic gain measured in Experiment 1 and 3, plotted as a function of target distance, and split according to the eccentricity range of the experimental session. Dots indicate the average gain, while the lines are the predictions of the multilevel model fit to the data. For the two conditions with smaller uncertainties (leftmost subpanels), average saccadic gains toward the intermediate targets (present in both 'large' and 'small' sessions) are overlapping, indicating that saccades were not systematically influenced by the eccentricity range of the targets. Only in the condition with the largest uncertainty (rightmost panel) did we find an effect of eccentricity range (i.e. a central bias). B. Size of the central bias, quantified as the parameter α of the regression model (see Results) and plotted as a function of the space constant (Experiment 1) or the peak luminance (Experiment 3) of the target. All error bars and bands are standard errors.

overshoot errors were considered about 7.5 times costlier (median across participants) than undershoots in Experiment 1, 95% CI [3.0, 15.7]; 6.7 times costlier in Experiment 2, 95% CI [2.9, 8.5] and 7.7 times costlier in Experiment 3, 95% CI [4.5, 18.4]. There was no significant difference in the estimated cost asymmetry across experiments, $F(2, 56) = 0.67, p = 0.51$. Overall, the assumption of an asymmetric, quadratic cost functions provides a good fit to variations in saccadic gain across all our experiments (Fig. 4). As an additional test, we used a leave-one-out cross-validation procedure to evaluate the predictive ability of the quadratic-asymmetric model against a descriptive model, which only assumed that the undershoot bias has a linear relationship with saccadic variability, without requiring that this relationship be adequate for minimizing an asymmetrical cost function. Across the three experiments, this test confirmed that assuming an asymmetric cost function results in a better and more parsimonious description of the data (see SI).

As an additional test of our hypothesis, we investigated whether gain variability could account for differences in gain, after controlling for the effects of our manipulations. For each experiment, we fit a multilevel linear model with the saccadic gain as dependent variable and luminance or space constant as categorical predictor, and participant as grouping factor. We took the residuals of these models and computed the correlation to the standard deviation of saccadic gain. We found a significant correlation (Pearson's r=-0.14, 95%CI [-0.26, -0.11]), which indicates that even after controlling for the influence of our manipulation, saccadic variability retains information about saccadic gain, a remarkable result given the individual differences in the degree of asymmetry of the cost function (see next section).

Cost asymmetry is related to the programming of corrective saccades. We examined further whether individual differences in the asymmetry of the cost function could be related to differences in the post-saccadic processing of the target. Across our three experiments we recorded a large number of secondary saccades (see SI for details), which can be appropriately defined as corrective because their amplitude was negatively correlated with residual error of the primary saccade (Fig. 5A). As mentioned in the Introduction, corrective saccades tend to have longer latencies when they are made in the direction opposite to that of the primary saccade (5–7), suggesting that overshoots and undershoots have different consequences for post-saccadic oculomotor processing. The latencies of small saccades, however, are also modulated by their amplitudes, which are often larger after undershoot errors (because they are larger, on average, than overshoots). To control for this effect, prior to segregating forward and backward corrective saccades (that is, in the opposite and same direction of the primary one, respectively), we fit a quadratic model to the latency of secondary saccades (as the dependent variable) as a function of their amplitudes (see SI and Fig. 5). We took the residuals of this model and classified them into forward and return saccades depending on the direction relative to the primary saccade. We then took, for each participant, the difference between the mean residuals of return saccades (which were expected to have longer latencies) and of forward saccades. This difference represents an estimate of the additional time cost required to prepare corrective saccades in the opposite direction to the primary one (Fig. 5B). Overall, this additional time cost was estimated to be about 30 ms, 95% CI [18, 44].

If the cost-function asymmetry that we estimated from the bias-variability relationship of primary saccades were related to this latency cost, then we should find a positive correlation between these two measures. Our data support this conjecture, providing clear evidence for a positive relationship (see Fig. 5), Pearson's r=0.50, 95%CI [0.28, 0.68]*. Computed separately for each experiment, the correlation estimates were: Experiment 1: r=0.60, 95%CI [0.04, 0.89]; Experiment 2: r=0.62, 95%CI [0.25, 0.84]; Experiment 3: r=0.46, 95%CI [0.07, 0.73]. To summarize, the joint analysis of secondary saccade latencies and primary saccade bias and variability indicates that the slower a participant is in correcting an overshoot error (relative to an undershoot), the more hypometric her/his saccades become with uncertainty about target location. This finding supports the notion that undershoots result from the visual system's strategy for keeping saccadic targets in the same visual hemifield (15), and extends that notion by showing that the parameters of primary saccades are optimized, taking into account the possibility that a secondary, corrective movement will be necessary.

Discussion

In the present study, we manipulated the positional uncertainty of a peripheral visual target and examined how the oculomotor system responded to increased uncertainty when planning saccades. In Experiment 1, we found that increasing the blur of the target (a Gaussian blob embedded in noise) produced a larger spread of the saccadic landing positions and decreased the precision of positional judgments in a related perceptual task. Crucially, as the uncertainty increased,

* To estimate the correlation we removed 3 data points (out of 59) corresponding to participants for which the standard error of the latency cost was larger than 30 ms (their mean standard error was ≈ 65 ms, whereas it was only ≈ 18 ms for the remaining participants). Adding these less reliable data points does not change the conclusions and yields a correlation of r=0.41, 95%CI [0.18, 0.60].

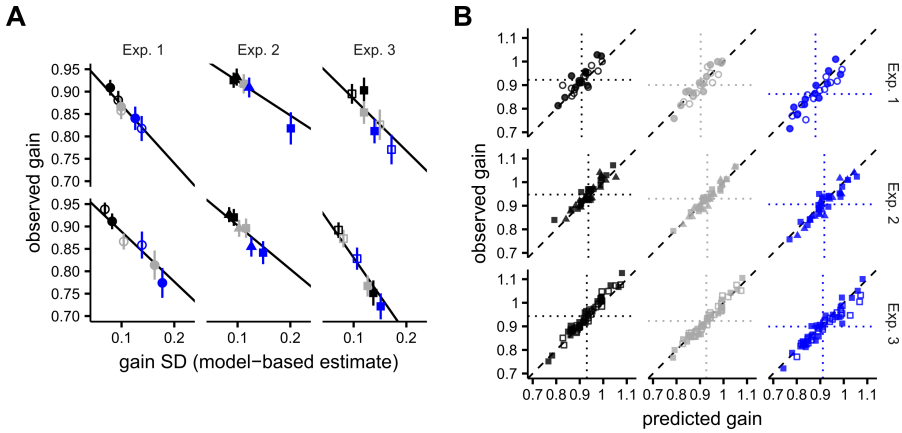

Fig. 4. Cost asymmetry determines the relationship between saccadic variability and bias. A. Estimated relationship between saccadic variability and bias for some example participants (two participants for each of the three experiments). The average saccadic gain for each condition and session is plotted as a function of the variability, as estimated by the model. Black lines represent the predicted gain, assuming the optimization of an asymmetrical, quadratic cost function. (Symbols follow the same conventions of Fig. 2, with the addition that for experiment divided in session with different eccentricity ranges filled and empty symbols indicate 'small' and 'large' sessions, respectively.) Error bars are 95% confidence intervals. B. Predicted and observed saccadic gain for all the participants, split by condition and experiment. The vertical and horizontal dotted lines indicate group means. See Fig. S1 for a similar plot showing observed and predicted standard deviations of saccadic gain.

saccades also became more hypometric, and systematically shifted toward the mean location of the target, a form of central tendency bias (13). The decrease in saccadic amplitude was well described by a simple model based on the assumption that the system is adapted to optimize an asymmetrical, quadratic cost function. In support of this assumption, we found that the estimated degree of asymmetry of the cost function was related across participants to the additional time required to plan a backward corrective saccade, made in the opposite direction to the primary one, relative to a forward one made in the same direction. In other words, the more time participants required to correct an overshoot (relative to an undershoot) with a secondary saccade, the more they decreased the mean amplitude of their primary saccades as the uncertainty in the target's position increased. These findings were corroborated by the results of Experiment 2 and 3, which also revealed that the reduced visibility of the target is the main source of these effects, while increasing the size of the target only produces a moderate, eccentricity-invariant decrease in saccadic amplitudes. Overall, the results presented here provide the first empirical evidence for theories arguing that an asymmetrical cost function is the source of the typical saccadic undershoot (15, 16), and establish experimentally the presence of a probabilistic mechanism that takes into account sensory and motor uncertainty to adjust where saccades are directed.

There are several (not necessarily incompatible) reasons for why the saccadic system might have evolved to avoid overshoot errors. According to one hypothesis (16), the system might seek to minimize the overall saccadic flight time: since visual sensitivity is much reduced during a saccade (17), it seems reasonable that the visual system may be adapted to maximize periods of clear view (even though the advantage would be only few milliseconds per saccade). Yet another hypothesis was advanced by Robinson (15), who proposed that the system may seek to maintain the post-saccadic target in the same visual hemifield as the pre-saccadic one, in order to facilitate further processing. This idea has been further developed by Ohl and colleagues (6, 18), who showed that secondary saccades are faster and more frequent after undershoots. These findings were interpreted in the context of a conceptual model, originally developed to explain the generation of microsaccades (19), which postulates that saccadic amplitudes are coded in a motor map endowed with short-range excitatory and long-range inhibitory connections. As a result, after each saccade the spatial distribution of neural activity would be biased toward the retinal location of the target in a way that facilitates further movements along similar direction, while slowing down movements in the opposite direction. If this imbalance represented an implementation constraint of the eye plant, then the system should take it into account by adopting a strategy that reduces the likelihood of overshoot errors. Therefore, Ohl's conceptual model (6, 18) provides a biologically plausible implementation of the cost function in our model, which was formulated at a more abstract, computational level of description. Our results support this conjecture, by showing that individual differences in the latency cost (see Fig 5C) are positively correlated with the estimated asymmetry of the cost function. Furthermore, additional analyses confirmed that individual differences in the latency cost were due to the difficulty in quickly planning backward corrective movements (see Supplementary Information, Fig. S3), rather than to the facilitation of forward corrections. This latter finding supports our interpretation that the functional role of saccadic hypometria is to avoid the slower corrections entailed by overshoot errors.

The present results help resolve a debate in the literature about the presence of a range effect (a central tendency bias) in saccadic targeting (9–12) by demonstrating that, although the range effect is not generally present when the target can be located with good precision, it does emerge when the positional uncertainty is large. In agreement with previous reports that 'averaging' saccades, which tend to fall in between the target and a distractor, are biased toward the most probable location of the target (20) our results support the view that a Bayesian process is working to optimize saccadic eye movement by taking advantage of prior knowledge. Although previous research suggested the saccades are normally based only on the most recent sensory information available (21–23), our current results show that when uncertainty is particularly high the saccadic system can reflect expectations developed over longer timescales, spanning multiple trials.

Finally, given that our experiment involved conditions of artificially high uncertainty that are uncommon in everyday life, one important issue in their interpretation is to what extent they generalize to more ecological conditions. While our experimental conditions were specifically designed to allow precise measurements of saccadic bias and variability under

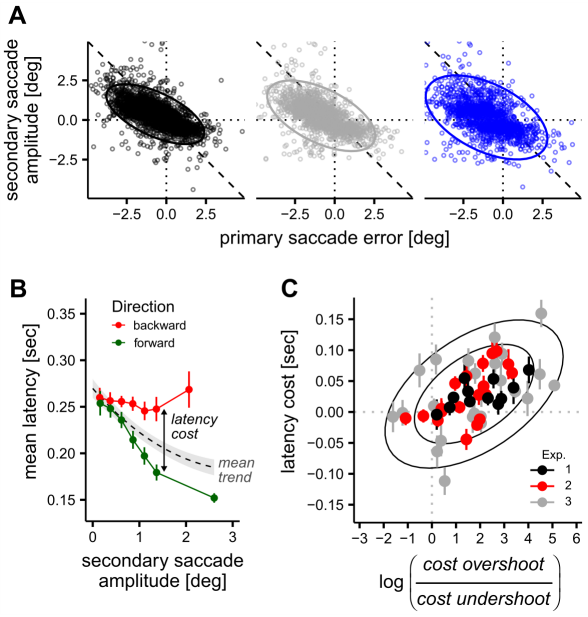

Fig. 5. Cost asymmetry is related to the programming of corrective saccades. A. Secondary saccades recorded in our experiments were corrective, as indicated by their negative correlation with the error of primary saccades. (Ellipses are 95% bivariate confidence intervals of the mean.) B. The latency cost is defined as the difference in latency between backward and forward saccades, after correcting for the mean trend due to the amplitude of secondary saccades (see SI for details, and Fig. S2 for a plot of saccadic latencies distributions). C. The relationship between estimated cost asymmetry (expressed as log-ratio of costs for an error of constant size) and the latency cost. See Fig. S3 in the Supplementary Information for separate analysis of the latencies of forward and backward saccades. Ellipses represent 75% and 95% bivariate confidence intervals.

conditions of varying uncertainty, previous studies have demonstrated that a systematic undershoot bias is present also under more ecological conditions, involving for example free viewing (24), visual search (25, 26), and free scanning of continuously present targets (27). High rates of error-correcting secondary saccades were found also under conditions designed to increase the difficulty of saccadic targeting during the scanning of stationary targets (28). In sum, the phenomena we examined in our study (saccadic undershoot and corrective saccades) are found also in a broad range of different and arguably more ecological experimental conditions, indicating that they reflect fundamental aspects of saccadic planning.

In conclusion, our results demonstrate that a flexible adaptive strategy underlies the control of saccadic amplitudes. By estimating the relationship between uncertainty about the target location, saccadic accuracy, and saccadic variability, we have shown that the typical undershoot bias of saccadic eye movements can be adequately explained as the result of strategy designed to optimize saccadic amplitudes, given sensorimotor uncertainty and an asymmetrical cost function. This strategy is probabilistic and Bayesian, in the sense that it must have at its disposal a trial-by-trial representation of uncertainty and it takes prior information into account. Together with previous reports that show how the distributions of saccadic landing positions are sensitive to rewards and task demands (29), the present results highlight the utility of eye-movement analysis as a tool to study probabilistic aspects of information processing in the brain.

Materials and Methods

See SI for the details of the experimental procedures and statistical analyses. All participants gave their informed consent in written form; the protocol of the study received full approval from the Research Ethics Committee of the School of Health Sciences of City, University of London. Data and code are available as an Open Science Framework repository: https://osf.io/293gc/.

ACKNOWLEDGMENTS. This work was supported by grant RPG-2016-124 from the Leverhulme Trust to M.J.M.

1. van Beers RJ (2008) Saccadic eye movements minimize the consequences of motor noise. *PloS one* 3(4):e2070.
2. van Beers RJ (2007) The sources of variability in saccadic eye movements. *Journal of Neuroscience* 27(33):8757–8770.
3. Harris CM, Wolpert DM (1998) Signal-dependent noise determine motor planning. *Nature* 394(6695):780–784.
4. Becker W, Fuchs aF (1969) Further properties of the human saccadic system: eye movements and correction saccades with and without visual fixation points. *Vision research* 9(10):1247–58.
5. Deubel H, Wolf W, Hauske G (1982) Corrective saccades: Effect of shifting the saccade goal. *Vision Research* 22(3):353–364.
6. Ohl S, Brandt SA, Kliegl R (2011) Secondary (micro-)saccades: The influence of primary saccade end point and target eccentricity on the process of postsaccadic fixation. *Vision research*.
7. Ohl S, Kliegl R (2016) Revealing the time course of signals influencing the generation of secondary saccades using Aalen's additive hazards model. *Vision Research* 124:52–58.
8. Vitu F, Casteau S, Adeli H, Zelinsky GJ, Castet E (2017) The magnification factor accounts for the greater hypometria and imprecision of larger saccades: Evidence from a parametric human-behavioral study. *Journal of Vision* 17(4):2.
9. Kapoula Z (1985) Evidence for a range effect in the saccadic system. *Vision Research* 25(8):1155–1157.
10. Kapoula Z, Robinson DA (1986) Saccadic undershoot is not inevitable: Saccades can be accurate. *Vision Research* 26(5):735–743.
11. Gillen C, Weiler J, Heath M (2013) Stimulus-driven saccades are characterized by an invariant undershooting bias: no evidence for a range effect. *Experimental Brain Research* 230(2):165–174.
12. Nuthmann A, Vitu F, Engbert R, Kliegl R (2016) No Evidence for a Saccadic Range Effect for Visually Guided and Memory-Guided Saccades in Simple Saccade-Targeting Tasks. *Plos One* 11(9):e0162449.
13. Hollingworth HL (1910) The Central Tendency of Judgment. *The Journal of Philosophy, Psychology and Scientific Methods* 7(17):461.
14. Watt RJ, Morgan MJ (1983) The recognition and representation of edge blur: Evidence for spatial primitives in human vision. *Vision Research* 23(12):1465–1477.
15. Robinson DA (1973) Models of the saccadic eye movement control system. *Kybernetik* 14(2):71–83.
16. Harris CM (1995) Does saccadic undershoot minimize saccadic flight-time? A Monte-Carlo study. *Vision Research* 35(5):691–701.
17. Holt EB (1903) Eye-movement and central anaesthesia. *The Psychological Review: Monograph Supplements* 4(1):1–45.
18. Ohl S, Rolfs M (2016) Saccadic Eye Movements Impose a Natural Bottleneck on Visual Short-Term Memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*
19. Rolfs M, Kliegl R, Engbert R (2008) Toward a model of microsaccade generation: the case of microsaccadic inhibition. *Journal of vision* 8(11):5.1–23.
20. He P, Kowler E (1989) The role of location probability in the programming of saccades: Implications for "center-of-gravity" tendencies. *Vision Research* 29(9):1165–1181.
21. Lisi M, Cavanagh P (2015) Dissociation between the Perceptual and Saccadic Localization of Moving Objects. *Current Biology* 25(19):2535–2540.
22. Massendari D, Lisi M, Cavanagh P, Collins T (2017) Is the efference copy of a saccade influenced by a perceptual illusion? *Journal of Vision* 17(10):879.
23. Lisi M, Cavanagh P (2017) Different spatial representations guide eye and hand movements. *Journal of Vision* 17(2):12.
24. Rasche C, Gegenfurtner KR (2010) Visual orienting in dynamic broadband (1/f) noise sequences. *Attention, Perception, & Psychophysics* 72(1):100–113.
25. McSorley E, Findlay JM (2003) Saccade target selection in visual search: Accuracy improves when more distractors are present. *Journal of Vision* 3(11):20.
26. FINDLAY JM (1997) Saccade Target Selection During Visual Search. *Vision Research* 37(5):617–631.
27. Findlay JM, Brown V (2006) Eye scanning of multi-element displays: I. Scanpath planning. *Vision Research* 46(1-2):179–195.
28. Wu CC, Kwon OS, Kowler E (2010) Fitts's Law and speed/accuracy trade-offs during sequences of saccades: Implications for strategies of saccadic planning. *Vision Research* 50(21):2142–2157.
29. Schutz AC, Trommershauser J, Gegenfurtner KR (2012) Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences* 109(19):7547–7552.

Supporting Information Text

Material and methods

Participants. In total, 55 naïve participants and 3 authors participated in 3 experiments. 10 naïve participants and 2 authors participated in Experiment 1 (mean age 37 years, SD 11.8; 2 females). 18 naïve participants and 2 authors participated in Experiment 2 (mean age 36 years, SD 12.1; 5 females). 25 naïve participants and 1 author participated in Experiment 3 (mean age 30 years, SD 9.1; 14 females). All participants had normal or corrected-to-normal vision and gave their informed consent in written form (the protocol of the study received full approval of the local ethics committee). Naïve participants were compensated with £8 for each hour of experiment.

Apparatus. During both perceptual and saccadic tasks, right eye-gaze positions were recorded with an Eyelink 1000 (SR Research Ltd., Mississauga, Ontario, Canada). The participant's head was placed on a chinrest with an adjustable forehead rest. Visual stimuli were presented on a gamma-linearized LCD monitor, 51.5cm wide, placed at 77cm of viewing distance. The monitor resolution was 1920×1200. An Apple computer controlled stimulus presentations and response collection; the experimental protocol was implemented using MATLAB (The MathWorks Inc., Natick, Massachusetts, USA), the PsychToolbox (1, 2) and the Eyelink toolbox (3).

Stimuli. Stimuli were Gaussian blobs presented on a background made of squares (side≈0.08 deg), with random luminance drawn from a Gaussian distribution (RMS contrast≈10%). In Experiment 1, blobs with different space constants (σ) were designed to have the same total energy, a manipulation that has already been shown to influence uncertainty about target position (4, 5). The peak luminance of the blob with smallest σ in the set corresponds to the maximum luminance that can be reached by the display (147 cd/m^2). When the peak coincided with a bright pixel of noise its luminance was set at this ceiling level. Three levels of σ were used: 0.3, 0.9 and 1.5 deg, which resulted in peak luminance values of 147, 57 and 50 cd/m^2, respectively. Experiment 2 was composed of two different conditions: in the first condition (fixed-size) σ was kept fixed at 0.9 deg, and we varied the peak luminance so as to match the peak luminance values obtained in Experiment 1. In the second condition (fixed-luminance), the peak luminance was fixed at the maximum value while σ varied in the same three levels of Experiment 1. Additionally, only in the fixed-luminance condition, we set the background noise to 0% RMS contrast. Experiment 3 used blobs with the same parameters as the fixed-size condition of Experiment 2.

We note that previous studies using brief, masked saccadic targets (6) found that both the precision of position judgments and the saccadic gain decreased with the duration of target presentation. However, a more recent study (7) examined in detail the effect of masks on saccadic programming, and concluded that a mask influences saccadic programming in the same way that a remote distractor does (8). The effect of the mask on saccadic amplitude depended in a complex way on target duration and characteristics of the visual mask (7). In particular, a tendency for the opposite bias (overshoot or hypermetria) was found for short presentation durations, when the mask was limited to the hemifield of the target rather than covering all of the monitor's width, suggesting that the effects of the mask are, at least in part, due to the spatial averaging (9) of mask and target. For these reasons, in the present study, we carefully avoided the use of spatially limited masks. One possible observation to our manipulations is that the more uncertain targets were also less salient. Although salience is an ill-defined concept, it is usually identified with low-level features such as luminance and contrast, which are known to bias saccadic landing positions when displays contain complex or multiple stimuli (10, 11). There is, however, little prior evidence that salience can influence saccadic accuracy in the case of single, simple targets such as our Gaussian blobs. For example, one study found saccadic accuracy to be roughly constant with respect to the luminance of the target, despite a large modulation of saccadic latencies (12).

Procedure. In all tasks, each trial started when gaze position was maintained within 2 deg from the central fixation point for at least 200 ms. If the trial did not start within 2 seconds, the program paused, allowing participants to take a break and re-calibrate the eye-tracker. Participants were encouraged to take a break whenever they felt the need to rest. To prevent the use of monitor edges as stable landmarks for the localization of the peripheral targets, the position of the fixation point was jittered across trials: each trial a new position was drawn from a circular, 2D Gaussian distribution centered on the screen center, with a standard deviation of 0.2 deg. The position of the peripheral targets (the Gaussian blobs) was always clamped with respect to the trial-by-trial position of the fixation point.

Perceptual task. The noise background (randomly generated each trial) appeared immediately after fixation was detected in the central position, followed after a random interval uniformly distributed within 300-500 ms by the two Gaussian blobs. The two blobs were placed at different distances on the left and right side of the fixation point. They were displayed for 250 ms, and then they disappeared together with the fixation point; the noise background instead remained visible until participants provided a response by pressing on the left/right arrow keys. The average distance of the two blobs was always 10 deg, while the difference in distance was adaptively adjusted using a separate QUEST+ (13) staircase for each blob's space constant. The procedure allowed us to select, for each trial and separately for each blob's σ, the stimulus that minimized the expected entropy of the three-dimensional posterior probability density of the parameter estimates of the psychometric function (a cumulative Gaussian psychometric function with symmetric lapse rate, i.e.

$$\Psi\left(x, \mu, \sigma_\psi, \lambda\right) = \lambda + (1 - 2\lambda)\Phi\left(\frac{x - \mu}{\sigma_\psi}\right), \qquad [1]$$

where Φ is the CDF of the standard normal distribution, λ the lapse rate, and σ_ψ is the shape parameter of the psychometric function, which will be referred to as JND). This method allows for the possibility that participants may have attentional lapses, while at the same time selecting ideal stimuli to constrain the parameters of the psychometric function. Participants performed one session of the perceptual task (12 blocks of 25 trials each) on each day of testing (in total each participant run 600 trials).

Saccadic task. The background noise appeared after fixation was detected in the central position, and it was followed by the Gaussian blob after a random interval uniformly distributed within 300-500 ms. In Experiment 1 the blob was located to the left or to the right of the fixation point, at an eccentricity of either 8, 9, or 10 deg in the "small" eccentricity session, or at either 10, 11, or 12 deg in the "large" eccentricity session (naïve participants were not informed about the difference between the two sessions; for the two authors the order was selected randomly and they were not given explicit information about the range of target eccentricity). These two blocks were run on different days, with order counter-balanced across participants. In Experiment 2 all the 5 eccentricities of the target were presented within a single session. In Experiment 3, the range was extended to 8 different eccentricities, equally spaced between 6 and 13 deg, which were split into two sessions (large vs small, with the target at 9 and 10 deg present in both sessions). In all tasks, each trial started when gaze position was maintained within 2 deg from the central fixation point for at least 200 ms. If the trial did not start within 2 seconds, the program paused, allowing participants to take a break and re-calibrate the eye-tracker. Participants were encouraged to take a break whenever they felt the need to rest. To prevent the use of monitor edges as stable landmarks for the localization of the peripheral targets, the position of the fixation point was jittered across trials: each trial a new position was drawn from a circular 2D Gaussian distribution centered on the screen center, with a standard deviation of 0.2 deg. The position of the peripheral targets (the Gaussian blobs) was always clamped with respect to the trial-by-trial position of the fixation point. In Experiments 1 and 2 the blob was displayed for 500 ms, while in Experiment 3 the duration was increased to 800 ms, intended to allow more time for secondary, corrective saccades. In all cases participants were instructed to shift their gaze onto it with a single saccade, as quickly and as accurately as possible. Trials in which participants blinked or moved their gaze before the appearance of the target were aborted and repeated at the end of the block. In Experiment 1 each session comprised 288 trials divided in 6 blocks (in total each participant run 576 trials); one of the participants ran the experiment split in 4 smaller sessions. In Experiment 2, each participant ran 2 sessions of the task for each of the two conditions (fixed-size and fixed-luminance); each session comprised 12 blocks of 15 trials (in total, each participant ran 720 trials). In Experiment 3, each participant ran 2 sessions, each comprising 240 trials, divided in 12 blocks. In all experiments the order of the different sessions (large vs small; fixed-size vs fixed-luminance) was counterbalanced across participants.

Analysis

Statistical analyses were performed using the free, open-source software R (14). Unless stated otherwise, group level estimates are reported as mean ± standard error, computed across participants and weighted by the number of trials (which could show slight variations across observers due different proportions of excluded trials). In the case of multilevel model estimates we reported the population level, or fixed-effect estimate, ± its standard error. Confidence intervals were obtained by bootstrapping (10^3 replications) using the bias-corrected and accelerated (Bca) method (15). χ^2 statistics indicate likelihood ratio tests between a full model and a reduced model where the specified parameter was constrained to be zero.

Perceptual task. To analyze psychophysical performance in our perceptual task while keeping into account the possibility that some error responses may be due to attentional lapses (i.e. stimulus-independent errors), we fit our data with 3 psychometric models that make different assumptions about the occurrence of lapses. The first model assumes that the lapse probability is always zero; the second one allows for a non-zero lapse probability that is assumed constant across conditions with different values of σ; the third model allows for lapse probability to vary across conditions (see Analysis section for details). We compared these different psychometric models using the AIC (Akaike Information Criterion). We found that for 10 out of 12 participants the model with fixed lapse rate was better than the model with varying lapse rate; and that for 7 out of 12 subjects the best model was the simpler one with lapse rate constrained to 0. Since we are interested in measuring how the positional uncertainty varies across conditions rather than in deciding which model provides a better description of the data, we averaged the estimates of the three models, weighting them according to the Akaike weights of each model (16), and performed group level analyses on the averaged estimates. The estimates of JNDs were: σ=0.3deg, JND=0.79±0.15deg (mean ± standard error); σ=0.9deg, JND=0.95±0.14deg; σ=1.5deg, JND=1.74±0.35deg.

Pre-processing of gaze recordings. Saccadic onsets and offsets were detected offline using MATLAB and an algorithm based on two-dimensional eye velocity (17). More specifically, saccades were identified as outliers in the two-dimensional velocity distribution of each trial and were identified as the part of gaze recordings that exceeded the median velocity by 5 standard deviations for at least 8 ms. Once saccadic parameters were measured, further statistical analyses were made using the open source software R (14). In our analysis we considered the whole sequence of saccades and microsaccades produced since the detection of the initial fixation to the end of the trial. For each trial, we selected as the primary saccade the first saccade that started after the onset of the target, from within a circular area of 2.5 deg around the initial fixation point, ended outside of that circular area, and had an amplitude of at least 1 deg. We excluded trials where the primary saccade had a latency shorter than 100ms or longer than 600ms. Since we were interested in studying the whole distribution of saccadic amplitudes for a given target distance, rather than just the saccades of a pre-specified amplitude, we applied only a loose filter on saccadic landing locations, by excluding only those trials where the landing location was more than 3 standard deviations away from

the mean landing location (computed separately for each eccentricity and condition). In order to reduce the error due to imprecisions in the eye-tracker calibration, we took the difference between the coordinates of the central fixation point, and the mean of all initial saccadic positions (made by a participant in a particular session) and used it to correct the initial and final saccadic positions. Finally, since the gaze is typically not exactly on the fixation point when the saccade starts, we normalized saccadic amplitudes in order to remove the variability due to trial-by-trial fluctuations in the fixation position by computing $S_n = T \times \frac{S}{E}$ where S is the raw saccade amplitude (distance between initial and final position), E is the retinal error of the target (distance between saccade initial position and target position), T is the distance of the target from the fixation point, and S_n is the saccade's normalized ampitude. In this article all saccade amplitudes reported are normalized according to this procedure, unless stated otherwise. Since the position of the saccadic targets differed from the initial fixation points only in its horizontal coordinates, in our analysis we considered only the horizontal components of saccadic amplitudes. Moreover, the vertical landing positions did not show any systematic bias or relationship with the horizontal component of saccadic amplitudes.

For the analysis of the range effect in Experiment 1 only, saccadic amplitudes were adjusted to eliminate differences across in mean baseline gain across sessions run on separate days. This correction was done by multiplying saccadic amplitudes in the session number i (with i=1,2) by the factor $c_i = 1 + G - G_i$ where G_i is the mean saccadic gain for a given subject in session i for the condition with smallest uncertainty (σ=0.3 deg) and G is the mean saccadic gain for the same condition but averaged over sessions. The correction is computed on the basis of the condition with σ=0.3 deg, because based on prior studies (18, 19) we did not expect any range effect in that condition, and before applying it we verified that this was the case also in our dataset (see Results). This allowed us to estimate how the central bias changed with respect to a baseline where uncertainty was minimal. This correction was not applied in the analysis of Experiment 3, in which both sessions were run in the same day.

For analyses involving secondary-corrective saccades, we included in the analysis the first saccade after the primary one, with a latency of at least 30 ms from the offset primary saccade. Since this interval may include some voluntary saccades made by the participant to shift back their gaze toward the center of the screen, in anticipation of the next fixation target, we excluded secondary saccades that increased the error of the primary saccade by more than 2.5 deg.

Analysis of saccadic landing positions. Multilevel models used in the analysis of saccadic landing positions were fitted using the R package lme4 (20). In all cases all fixed-effects parameters had corresponding random effects, grouped according to the participant. A fully parametrized, random effects variance-covariance matrix was estimated in all cases.

Bayesian multilevel models (used in the analysis of the central bias) were estimated using Stan (21) and its R interface. In both Experiment 1 and Experiment 3, we fit the models using MCMC sampling to approximate the posterior distribution of the parameters. We ran six Markov chains of 2000 samples each, and verified convergence by checking that there were no divergent transitions and that the variance between and within chains did not differ significantly; the $\hat{R}$ statistic was smaller than 1.1 for all parameters (22). Beta coefficients were given weakly informative Gaussian priors, with standard deviation of 2, centered on zero for the intercepts and on 1 for the slopes of saccadic amplitudes. Compression parameters α were given Guassian priors centered on zero and with a standard deviation of 1. Bayesian credible intervals were obtained using the percentile method on the samples from the posterior distribution.

In the analysis of secondary saccades, the use of a quadratic model is motivated by the observation that the relationship between secondary saccade latency and amplitude was not perfectly linear. Latency decreased faster for smaller amplitudes, and then tended to asymptote toward a minimum, as the amplitude increased. To better summarize this relationship, we added a quadratic term to the model, so that the expected value of the saccadic latency was modelled as a second-degree polynomial function of the saccadic amplitude. We included also the experiment (1, 2, or 3), with interactions for both the linear and quadratic terms, giving the model a total of 9 free parameters: 9 for the fixed effects (the three coefficients of a second-order polynomial times the three experiments) and 6 for the random effects (the elements of the variance-covariance matrix of a trivariate normal distribution). The model was fit to 7885 secondary saccades.

Estimation of asymmetric cost function. It has long been known that saccades display not only a variable error but also a constant error or bias. This bias has been found also in other primate species (e.g. (23)), is typically larger in infants and gradually decreases during development (24) and some evidence suggests that it is present also in free-viewing tasks (25). While it is clear that the variable error originates from a combination of uncertainty in the estimated location of the target and motor noise (26), the origin of the undershoot has been long debated. Studies of oculomotor adaptation have provided the first empirical evidence that the undershoot bias may be a deliberate strategy: when the target position is moved during the saccade, so as to cancel the undershoot bias, the saccadic system quickly learns to undershoot the new, altered, postsaccadic location of the target (27). This empirical observation suggested that the oculomotor system is willing to tolerate a small undershoot bias in order to reduce the probability of overshoots. In the present study, we have formalized this idea by assuming that the system is optimizing an asymmetrical cost function that assigns a larger cost to overshoot errors than to undershoot of the same size. Specifically, we assumed that the saccadic system seeks to optimize the asymmetric cost function

$$L(x) = [a + (1 - 2a) \cdot 1_{x<0}] \cdot x^2 \qquad [2]$$

where x is the landing error in gain units (if S is the saccadic amplitude, and E the target eccentricity, then $x = S/E - 1$), $1_{x<0}$ is an indicator function (equal to 1 when the argument in the subscript is verified, and 0 otherwise) and a is the asymmetry parameter, bounded between 0 and 1. Assuming that the variability of saccadic gain is well approximated by a Gaussian

distribution, the expected cost for a given level of gain variability σ_g and cost asymmetry a can be expressed as

$$\begin{aligned}\mathbb{E}\left[L \mid \mu, \sigma_g, a\right] &= \int_{-\infty}^{\infty} L(x) \cdot \frac{1}{\sigma_g}\phi\left(\frac{x-\mu}{\sigma_g}\right) dx \\ &= \frac{\mu^2+\sigma_g^2}{2} + (1-2a)\left[\frac{(\mu^2+\sigma_g^2)\cdot \operatorname{erf}\left(\frac{-\mu}{\sqrt{2}\sigma_g}\right)}{2} - \frac{\mu\cdot\sigma_g\cdot e^{-\frac{\mu^2}{2\sigma_g^2}}}{\sqrt{2\pi}}\right]\end{aligned} \quad [3]$$

where ϕ is the probability density function of the standard normal distribution, $\phi(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2}$, and erf the error function, $\operatorname{erf}(x) = 2/\sqrt{\pi}\int_0^x e^{-t^2}dt$. Note that when $a = 0.5$ (i.e. the cost function is symmetric) the last term cancels out and the minimum of the expected cost is obtained for $\mu = 0$, that is when the saccadic system aims precisely at the center of the target (by trying to produce a saccade with gain equal to 1). The above equations allow finding the ideal gain, $\arg\min_\mu \mathbb{E}\left[L \mid \mu, \sigma_g, a\right]$, that is the aimpoint μ that yields the minimum of the expected cost for arbitrary levels of saccadic variability σ_g and cost asymmetry a. The ideal gain and the variability level can be used to specify a likelihood function for the observed distribution of saccadic errors for a given condition, i.e.

$$p\left(x \mid \sigma_g, a\right) = \frac{1}{\sigma_g}\phi\left(\frac{x - \arg\min_\mu \mathbb{E}\left[L \mid \mu, \sigma_g, a\right]}{\sigma_g^2}\right). \quad [4]$$

This likelihood function can be used to estimate the value of the parameters that maximize the probability of the observed data. We used Brent's method (28) to find the ideal gain for given combinations of parameters, and we used box-constrained optimization (with $0 < a < 1$ and all $\sigma_g > 0$), as implemented in the optim() function in R and the L-BFGS-B (29, 30) algorithm, to identify the parameters values that maximized the likelihood of the observed distributions of saccadic errors. The median of maximum-likelihood estimates of the asymmetry parameter (together with their 95% confidence intervals) was 0.87 [0.75, 0.94] in Experiment 1, 0.87 [0.74, 0.89] in Experiment 2, and 0.88 [0.82, 0.95] in Experiment 3. Individual differences in the degree of asymmetry of the cost function were examined by analyzing the logarithm of the cost ratio between overshoots and undershoots (given a fixed error magnitude), which can be calculated as $\log\left(\frac{\alpha}{1-\alpha}\right)$. We preferred to use the logarithm of the ratio in order to compute the correlations reported in the main text, because its distribution is not different from normal according to a Kolmogorov-Smirnoff test, D=0.07, p=0.92, whereas both the distributions of the parameter α and that of the simple ratio did deviate significantly from normality:α, D=0.20, p=0.01; cost-ratio, D=0.28, p=0.0001.

To assess the predictive ability of our model based on the quadratic-asymmetric cost function, we compared it against an alternative, descriptive model by means of a cross-validation test. The alternative, descriptive model assumed only that the undershoot bias has a linear relationship with saccadic variability, without requiring that this relationship be adequate for minimizing an asymmetrical cost function. More specifically, while the relationship in the asymmetric-cost model is specified by a single parameter (the parameter which determines the asymmetry of the function), in the null model this relationship is determined by two parameters, an intercept (β_0) and a slope (β_1). The likelihood function for this model can be expressed as

$$p\left(x \mid \sigma_g, \beta_0, \beta_1\right) = \frac{1}{\sigma_g}\phi\left(\frac{x - \beta_0 - \beta_1\sigma_g}{\sigma_g^2}\right). \quad [5]$$

In order to quantitatively evaluate the predictive ability of the quadratic-asymmetric model, we performed a cross-validation test. For each participant we iteratively estimated the model, keeping the data from one condition aside as test set. In each experiment there were six of such conditions (three levels of uncertainty times 2 sessions). In the cross-validation test we iteratively fitted the model on five of these, and used the estimated parameters to predict the hold-out test condition. Overall, across the three experiments, the cross-validated log likelihoods (summed over test set for each participant) indicated that for 43 out of 59 cases the quadratic-asymmetric model was better at predicting the test set than the null model. The mean log-likelihood difference (quadratic minus null) was 17853 (SD: 10164, range: -60 to 580115). To summarize, this result indicates that assuming an asymmetric cost function results in a better and more parsimonious description of the data.

Supplemental results

Perceptual precision decreases with increasing blur of the targets. In our first experiment positional uncertainty was manipulated by varying the space constant (σ) of a Gaussian blob embedded in noise. The maximum luminance of this saccadic target co-varied, such that its contrast energy remained constant. To verify that our manipulation did indeed affect positional uncertainty we used a perceptual bisection ta sk, in which two targets were simultaneously presented for 250ms and participants (n=12) were asked to indicate which of the two was the furthest from an intermediate fixation point. Our dependent variable was the just noticeable difference, or JND, quantified here as the reciprocal of the psychometric slope. A repeated-measures ANOVA confirmed that the JND increased with the space constant of the Gaussian blob, $F(2, 22) = 16.74$, $p = 4.48 \times 10^{-4}$. In sum, the results of the perceptual task confirmed that our manipulation successfully affected the positional uncertainty of the target, presumably because of imperfect spatial integration of target's luminance contrast, which would make these larger targets harder to see.

Relationship between saccadic latency and undershoot. The undershoot also varied with saccadic latency, with a tendency toward greater undershoots with longer latencies. This finding precludes any speed-accuracy trade-off (cf. (31)). We binned trials according to the quartiles of individual latency distributions, and run a two-way repeated measures ANOVA to assess whether saccadic gain varied as a function of latency. We found an interaction between latency quartile and σ, suggesting that the effect of latency tended to become larger as σ increased [latency quartile, $F(3,33) = 2.40, p = 0.086$; σ, $F(2,22) = 16.21$, $p = 4.71 \times 10^{-5}$; interaction: $F(6,66) = 2.68, p = 0.022$]. This pattern is opposite to what has been reported in a study of target displacement (i.e., from fixation to the parafoveal visual field; (32)), wherein saccadic delay allows a low-pass filtered signal of the target's position to approach its asymptote. In our case, fluctuations in saccadic latency were likely caused by differences across trials in the random noise background and its effect on target visibility. Indeed, less visible targets, which resulted in longer latencies, produced also a greater spread of saccadic landing positions [latency quartile, $F(3,33) = 5.09, p = 0.005$; σ, $F(2,22) = 7.41, p = 0.003$; interaction: $F(6,66) = 1.55, p = 0.174$], suggesting that target visibility is the main modulator of positional uncertainty.

References

1. Brainard DH (1997) The Psychophysics Toolbox. *Spatial Vision* 10(4):433–436.
2. Kleiner M, Brainard D, Pelli D (2007) What's new in Psychtoolbox-3. *Perception* 36(14):1–16.
3. Cornelissen FW, Peters EM, Palmer J (2002) The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior research methods, instruments, & computers : a journal of the Psychonomic Society, Inc* 34(4):613–617.
4. Izawa J, Shadmehr R (2008) On-line processing of uncertain information in visuomotor control. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 28(44):11360–8.
5. Watt RJ, Morgan MJ (1985) A theory of the primitive spatial code in human vision. *Vision Research* 25(11):1661–1674.
6. Aitsebaomo AP, Bedell HE (1992) Psychophysical and saccadic information about direction for briefly presented visual targets. *Vision Research* 32(9):1729–1737.
7. Guillaume A (2012) Saccadic inhibition is accompanied by large and complex amplitude modulations when induced by visual backward masking. *Journal of Vision* 12(6):5–5.
8. Edelman JA, Xu KZ (2009) Inhibition of Voluntary Saccadic Eye Movement Commands by Abrupt Visual Onsets. *Journal of Neurophysiology* 101(3):1222–1234.
9. Findlay JM (1982) Global visual processing for saccadic eye movements. *Vision Research* 22(8):1033–1045.
10. Schutz AC, Trommershauser J, Gegenfurtner KR (2012) Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences* 109(19):7547–7552.
11. Stritzke M, Trommershäuser J, Gegenfurtner KR (2009) Effects of salience and reward information during saccadic decisions under risk. *Journal of the Optical Society of America* 26(11):B1–13.
12. Doma H, Hallett PE (1988) Dependence of saccadic eye-movements on stimulus luminance, and an effect of task. *Vision Research* 28(8):915–924.
13. Watson AB (2017) QUEST+: A general multidimensional Bayesian adaptive psychometric method. *Journal of Vision* 17(3):10.
14. R Core Team (2015) R: A language and environment for statistical computing.
15. Efron B (1987) Better Bootstrap Confidence Intervals. *Journal of the American Statistical Association* 82(397):171.
16. Burnham KP, Anderson DR (2002) *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach.* (Springer New York, New York, US), 2nd editio edition, p. 488.
17. Engbert R, Mergenthaler K (2006) Microsaccades are triggered by low retinal image slip. *Proceedings of the National Academy of Sciences of the United States of America* 103(18):7192–7.
18. Gillen C, Weiler J, Heath M (2013) Stimulus-driven saccades are characterized by an invariant undershooting bias: no evidence for a range effect. *Experimental Brain Research* 230(2):165–174.
19. Nuthmann A, Vitu F, Engbert R, Kliegl R (2016) No Evidence for a Saccadic Range Effect for Visually Guided and Memory-Guided Saccades in Simple Saccade-Targeting Tasks. *Plos One* 11(9):e0162449.
20. Bates D, Mächler M, Bolker B, Walker S (2012) Fitting linear mixed-effects models using lme4. *Journal of Statistical Software ...* 67(1):51.
21. Carpenter B, et al. (2017) Stan : A Probabilistic Programming Language. *Journal of Statistical Software* 76(1).
22. Gelman A, Rubin DB (1992) Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science* 7(4):457–472.
23. Opris I, Barborica A, Ferrera VP (2003) Comparison of performance on memory-guided saccade and delayed spatial match-to-sample tasks in monkeys. *Vision Research* 43(3):321–332.
24. Harris CM, Jacobs M, Shawkat F, Taylor D (1993) The development of saccadic accuracy in the first seven months. *Clinical Vision Science* 8(1):85–96.
25. Rasche C, Gegenfurtner KR (2010) Visual orienting in dynamic broadband (1/f) noise sequences. *Attention, Perception, & Psychophysics* 72(1):100–113.
26. van Beers RJ (2007) The sources of variability in saccadic eye movements. *Journal of Neuroscience* 27(33):8757–8770.
27. Henson DB (1978) Corrective saccades: Effects of altering visual feedback. *Vision Research* 18(1):63–67.
28. Brent R (1973) Chapter 4: An Algorithm with Guaranteed Convergence for Finding a Zero of a Function in *Algorithms for Minimization without Derivatives.* (Prentice-Hall, Englewood Cliffs, NJ).

29. Byrd RH, Lu P, Nocedal J, Zhu C (1995) A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing* 16(5):1190–1208.
30. Zhu C, Byrd RH, Lu P, Nocedal J (1997) Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software* 23(4):550–560.
31. Coëffé C, O'regan JK (1987) Reducing the influence of non-target stimuli on saccade accuracy: Predictability and latency effects. *Vision Research* 27(2):227–240.
32. de Bie J, van den Brink G, van Sonderen J (1987) The systematic undershoot of saccades: a localization or an oculomotor phenomenon? in *Eye Movements from Physiology to Cognition*. (Elsevier), pp. 85–94.

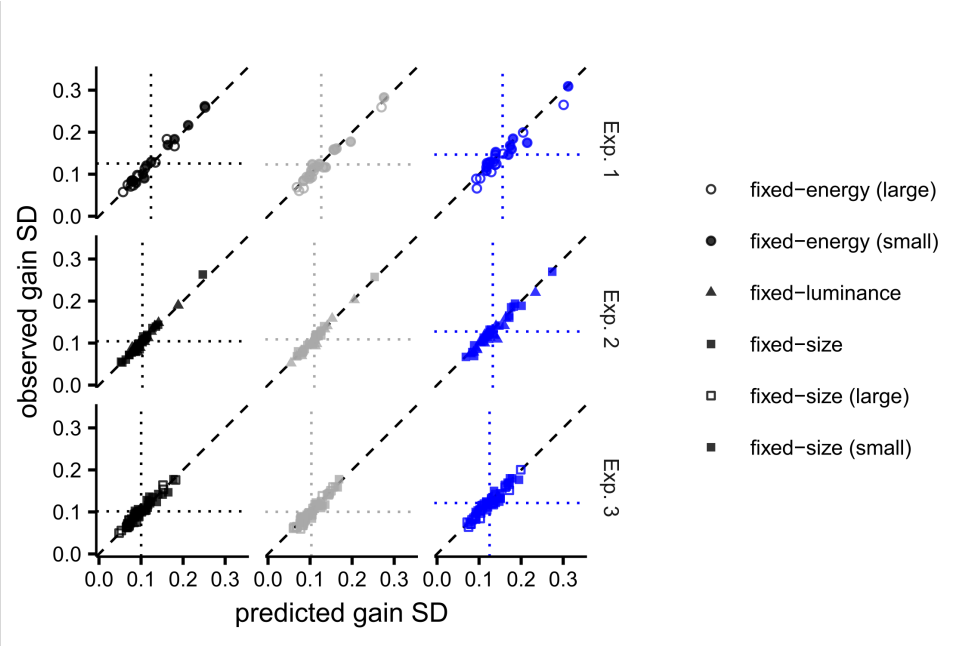

Fig. S1. Relationship between observed and predicted (that is the model-based estimate σ_g) standard deviation of saccadic gain. Conventions are the same as Fig. 4B, Main text.

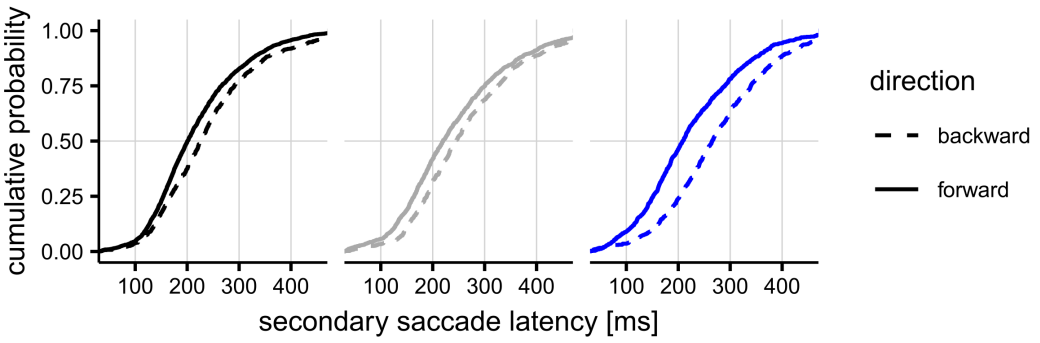

Fig. S2. Empirical distribution functions of secondary saccadic latencies, pooled across all experiments and divided according to the direction (forward vs. backward, continuous and dashed lines, respectively) and the conditions of expected uncertainty (from right to left, increasing blur, decreasing luminance or increasing size).

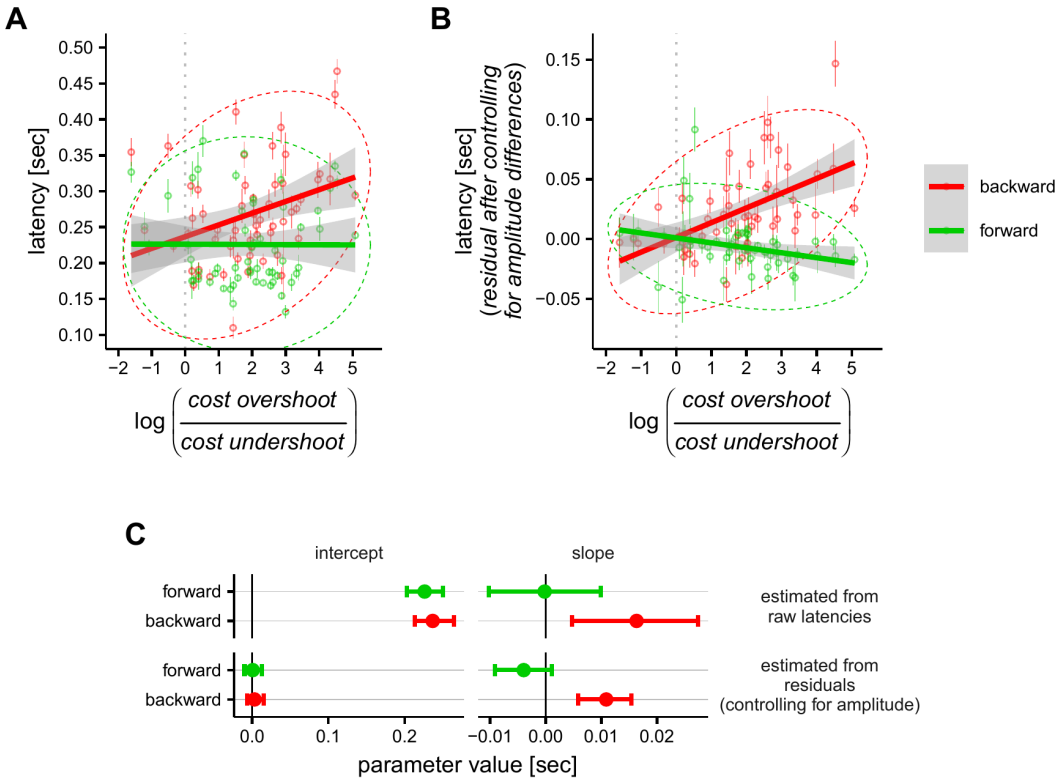

Fig. S3. Additional analyses of secondary saccade latencies. To further investigate the origins of the individual differences in the latency cost that we measured (see Fig. 5B and 5C in the main text), we examined the relationship between the estimated cost asymmetry and the latency of forward and backward secondary saccades. Panel A represents the raw latencies, split according to direction (forward vs backward), and plotted as a function of the log cost asymmetry. Ellipses represents 95% bivariate confidence intervals of the mean, and the lines show linear regressions with 95% confidence bands. Panel B represents the same analysis but performed on the *residual* individual differences in latencies, that is after removal of the mean effect of saccadic amplitude (as shown in Fig. 5C). The values of the parameters, together with bootstrapped 95% CI are shown in panel C. The intercept parameter represents latency when the estimated ratio of undershoot and overshoot costs is 1 (and therefore the log cost ratio is equal to 0, indicating a symmetrical cost function). It can be seen that this parameter does not differ significantly between forward and backward saccades. This indicates that, for subjects with very small asymmetry (log cost-ratio ≈ 0), the latency of secondary saccades was similar regardless of the direction (forward vs backward). However, as the estimated cost asymmetry increases, we find that the average latency of secondary backward saccades increases systematically (see the slope parameters in panel C), whereas that of the forward saccades remains constant. In other words, the between-subjects variability in the latency cost that we measured (i.e. the latency of backward saccades minus that of forward saccades) is due to backward saccades being slower in subjects with greater cost asymmetry, and not from forward saccades being faster. To summarize, this analysis reveals that subjects with a greater cost asymmetry (i.e. those who, given a certain increase in saccadic variability, display the largest increase in the undershoot bias) are characterized by comparatively slower backward secondary saccades; not faster forward ones. Thus, the relationship between saccadic variability and undershoot bias appears to be determined by the increased difficulty (for the oculomotor system) in programming a backward corrective saccade. This finding therefore supports our hypothesis that participants who have slower backward saccades avoid costlier overshoot errors by undershooting more, as end-point variability increases.