Supporting Information
Identifiability for the One-Choice Diffusion Model

There is an identifiability issue in the one-choice model. The essence of the issue is that boundary setting, drift rate, and standard deviation (SD) in drift rate are highly correlated. For these parameters, different starting values for the program that fits the model to data can yield different best-fitting parameter estimates such that the different estimates have about the same goodness of fit. The differences in the best-fitting parameter estimates can be as large as 2:1. The correlations between the parameters are over .9. In a 3-dimensional representation, the three parameters covary on a single straight line. However, ratios of the parameters are almost invariant over the different fits. The ratio of drift rate to SD in drift rate (v/η) and the ratio of boundary setting to drift rate (a/v) are almost constant.

One way to understand the identifiability issue is to consider what the model must explain. For one-choice RT tasks, the only data are the shapes and locations of the RT distributions. In contrast, for two-choice tasks, the data are made up of the RT distributions for correct and error responses as well as accuracy. This larger number of constraints makes the two-choice model identifiable (Ratcliff & Tuerlinckx, 2002). For the identifiability issue with the one-choice model, the parameters trade off in a simple way: If the boundary setting is increased, then responses slow. To compensate for this, drift rate and across-trial variability in drift rate (and to a lesser extent, nondecision time) are increased.

In order to examine the identifiability issue, we provide fits of the model to data using different starting values and describe the results of a series of simulations.

Starting Values

We fit the data for the two conditions of Experiment 2 (sleep-deprived and non-deprived) with three different sets of starting values for a, v, and η. The sets were chosen so as to produce large differences in the best-fitting values of the parameters. Table S1 shows that they did so. In the table, the first, second, and third rows for the two conditions correspond to the three different sets of starting values. The best-fitting parameter values differ by about 50% from smallest to largest, yet the chi-square goodness of fit measure does not differ significantly.

Figure S1 shows how similar the predicted RT distributions are for the three sets of starting values. The top panel shows smoothed frequency polygons (histograms, based on 20,000 simulated RTs) and the bottom panel shows the corresponding hazard functions (based on the muhaz function in R which uses an Epanechnikov kernel method). The range used for the hazard function estimation was the .001 quantile RT to the .97 quantile RT. The deviations between the three frequency polygons and between the three hazard functions are small and show no consistency. In fact, the deviations between the three functions for sleep-deprived and nondeprived conditions are no greater than the variability across random samples.

The inability to obtain unique parameter estimates is a problem because it means that it is not possible to uniquely estimate these components of processing. However, the three sets of parameter values provide two invariant ratios of any of the pairs of the three parameters. Even though the smallest values of the parameters for boundary setting, drift rate, and SD in drift rate are about 2/3 the maximum value, the ratios are constant to within about 5%. This means that the model does extract invariances: invariances that are the ratios of parameter values.

Simulations

We designed simulations to illustrate the invariance of the parameter ratios. In the first set of simulations, we fit the model to 30 sets of data simulated from the model with the same parameter values for each data set and 2,000 observations per distribution (similar to Experiment 2). The model was fit to the simulated data sets with randomly generated starting values designed to produce a range of parameter values like those in Table 1. This produced differences in parameter values for drift rate, SD in drift rate and boundary separation as large as 2:1 allowing correlations among parameter values to be obtained to determine whether the parameters trade off as in Table 1. Figure S2 shows results from two simulations with high and lower drift rates similar to those in the fits to the sleep-deprived and nondeprived conditions in Experiment 2. In both cases, the correlations between boundary separation, drift rate and SD in drift rate were very high. These parameters also fall on a straight line in 3 dimensions (visualized using the GGobi package from R).

The results showed a property of the model that provides ambiguity in interpretation. Although the ratios of boundary separation, drift rate and SD in drift rate are almost constant, the nondecision time is correlated with these parameters. This presents a problem with estimating nondecision time; the best that can be done is a range of values. However, the range of values is not large and the value of T_{er} can be determined to within plus or minus 20 ms. It can also be specified with reference to the other model parameters, for example, boundary separation. These have the highest correlations.

Figure S3 shows simulations using three conditions as in Experiment 1, with a set of parameter values used to generate simulated data that match those in Experiment 1. These show the same correlations as for the one-condition simulations in Figure S2. Therefore, multiple stimulus conditions do not help constrain the model parameters and the identifiability issue is the same with multiple conditions.

What Parameters are Necessary?

For the two-choice model, it is necessary to assume that the components of processing for the model vary from trial to trial (Ratcliff et al., 1999; Ratcliff & Tuerlinckx, 2002). Firstly, this is motivated by common sense, in that it is not plausible that subjects could accurately set decision parameters to the same values from trial to trial, nor is it plausible that the evidence from stimuli that is used in making a decision is identical from trial to trial even for nominally identical test items.

Secondly, parameter variability is necessary for the model to fit the relative patterns of error and correct RTs. Variability in drift rate across trials produces errors slower than correct responses, and variability in starting point across trials (equivalent to variability in boundary settings, for low to moderate values) produces errors faster than correct responses. These two together can produce patterns such as errors faster than correct responses in high-accuracy conditions, errors slower than correct responses in low-accuracy conditions, errors faster than correct responses in speed-instructed conditions, and errors slower than correct responses in accuracy-instructed conditions. Variability in nondecision time across trials can produce shifts in the leading edges of RT distributions between high- and lower-accuracy conditions. All of these sources of variability are identifiable in the two-choice model. When data are simulated from a set of parameter values, the fits of the model to the data recover the parameter values accurately. However, it should be noted that the model's variability parameters are less well estimated, as shown by the large SDs in their estimates. For a single session of data (e.g., Experiment 1 in the article), the SDs can be so large that modest differences in their values across subject groups are often not significant.

In the one-choice model, there are many fewer degrees of freedom in the data. To determine which variability parameters of the model are necessary to fit experimental data, we fit the model to the data from the two experiments in the body of the paper and the experiment described below, taking out parameters one at a time. From the chi-square goodness of fit values shown in Table S2, we can determine which parameters improve goodness of fit when introduced and which do not (a reduction of one degree of freedom will change the significance level by a little more than 1). Note that the fits are done by simulation so the goodness of fit values have variability as is evidenced by the turn around for Experiment 1 with three values of SD in drift, variability in the nondecision time versus the same case with variability in decision criterion.

Adding variability in the decision boundary (or equivalently the starting point) improves goodness of fit only marginally or not at all. Relative to the model with variability in drift rate and nondecision time, taking out variability in drift rate produces the largest chi-square values. Taking out variability in nondecision time produces the next largest chi-square values. The conclusion is that variability in drift rate and variability in nondecision time are necessary to accurately fit the model to data.

Are Negative Drift Rates Necessary for Good Fits?

We refit the model to the data from the sleep-deprived and nondeprived data sets (the data with the greatest number of observations per RT distribution) with the drift rate only allowed to be zero or positive. If the drift rate from the normal distribution of drift rates across trials was selected to be negative in the fitting program, it was set to zero. The bottom two rows in Table S1 shows the results. The fits are almost the same, with similar parameter values to the model with negative drift rates. Thus, if we restrict the model to apply only to processes that are detected, this is an option that allows all processes to have positive or zero drift and so they would all eventually terminate if enough time were allowed.

RT Histograms for Experiment 2

In the body of the paper, for Experiment 2, we focused on the fits of the model to hazard functions. We grouped data across similar performance levels to produce distributions with about 2,000 observations per subject per condition.

This provided reasonably good estimates of the hazard functions. The literature usually presents RT distributions, although they are not as decisive for evaluating the behavior of the tail of the distribution. For completeness, Figure S4 displays histograms for the RT data and for the RTs predicted from the best-fitting parameter estimates; these correspond to the hazard functions in Figure 3 in the body of the article.

Experiment S1: Adding a Variable Onset Two-Choice Condition

In Experiment 1 in the body of the article, stimuli were 20x20 arrays of black and white pixels centered in 100x100 arrays. The 20x20 arrays varied from one condition of the experiment to another in the number of pixels that were black. For the one-choice task, the delay between the onset of the 100x100 array and the onset of the 20x20 array was variable. For the two-choice task, it was not; the 20x20 and the 100x100 array were displayed simultaneously. It is a natural question then to ask what would happen in the two-choice task if the onset of the 20x20 array was variable.

To address this, we conducted an experiment in which we modified Experiment 1 to add variable-onset conditions for the two-choice task. The variable onset for the two-choice task was the same as for the one-choice task: the 20x20 stimulus was displayed after the 100x100 array with the same timing as for Experiment 1 (500 ms plus an exponentially distributed amount of time with mean 750 ms). Unlike in Experiment 1, we used the same levels of brightness for the one-choice and two-choice tasks. The proportions of black to white pixels or white to black pixels in the 20x20 stimulus arrays were .60 and .56 for all tasks (i.e., four brightness levels).

There were three tasks in the experiment: the one-choice task, the two-choice task without variability in stimulus onset time, and the two-choice task with variability in stimulus onset time. The time-out was 2,000 ms after the stimulus presentation. A blocked design was used with 64 trials per block and with the tasks precued at the beginnings of blocks. There were 19 subjects from the same population as Experiment 1. In all other details, the method was the same as for Experiment 1.

Results

Table S3 shows the results of the experiment for the three tasks and the two difficulty conditions (.60 and .56 pixel probability): accuracy for the two-choice tasks and, for all three tasks, mean RTs, the .1 quantile RTs (which represent the leading edges of the RT distributions), and the proportions of responses that timed out at the 2,000 ms limit. Table S4 shows the best-fitting values of the models' parameters.

The two-choice tasks.For the two types of two-choice tasks, accuracy for responses made under 2,000 ms is near ceiling (Table S3), varying from about .96 to about .99. With 96 trials per brightness condition and accuracy at .98, for example, the number of errors per subject would be only about two. In fact, any observed errors may be spurious, i.e., hitting the wrong response key by mistake. High accuracy is a consequence of the way the fixed-onset and variable-onset tasks must be designed. If the stimuli had been made more difficult, then the number of errors would have been larger for the fixed-onset task, making modeling for that task more constrained. However, for the variable-onset task, this change might have led to a large number of trials on which a subject did not detect the stimulus at all, whereas for the trials in which the stimulus was detected, accuracy would still have been high, because detecting the onset of the stimulus almost always entails knowing whether it was bright or dark. (In Experiment 1, which did not include the variable-onset two-choice task, the levels of brightness were set to make the task more difficult and accuracy values were not at ceiling.)

For the fixed-onset task, the difference between the .1 quantiles for the two difficulty conditions was small, only 23 ms, and the difference in the mean RTs was 84 ms. For the variable-onset task, the differences were much larger, 80 ms for the .1 quantile and 264 ms for mean RTs. In sum, the target for modeling was high accuracy in both tasks, with relatively small changes in RT distributions between the brightness conditions for the fixed-onset task but larger changes for the variable-onset task.

For both tasks, the two-choice diffusion model was fit to the group data, not to data for individual participants, because there were so few errors. Specifically, there were 22 out of 152 subject-conditions (2 tasks by 4 brightness levels by 19 subjects) with zero errors and 108 out of 152 with fewer than 5 errors. The data input to the fitting program were the same as for the Experiment 1 in the body of the article except for error RTs. Using five RT quantiles was not possible because of the small numbers of errors. Therefore, a median split was used, dividing the error RT distribution into halves.

If the diffusion model were to miss the 80 ms shift in the RT distribution for the variable-onset task, then there would be mis-predictions for the .1 quantiles. Another way to examine this issue is to allow the model to fit the data with different values of nondecision time for the two difficulty conditions. Then the difference between the two best-fitting values of nondecision time would be an estimate of the misfit. It was this latter method that we used.

The results show that, for both the fixed- and variable-onset tasks, the best-predicting values of nondecision time (Table S4) were about the same for the two difficulty conditions. For the fixed-onset task, the difference in the estimates for the two conditions was only 4 ms, and for the variable-onset task, the difference in the estimates was 25 ms (much less than the 80 ms difference in the .1 quantiles in the data). Thus, the two-choice model appears to fit the two-choice data for both fixed-onset and variable-onset conditions reasonably well. Note that because of the low numbers of errors, the model was fit to group data and so absolute goodness of fit measures are not provided.

There were modest differences in the boundary separation parameters between the two tasks (Table S4), probably because the variable-onset task was perceived to be more difficult than the fixed-onset task. Also, the SD in drift rate across trials (η) was quite different for the two tasks, but this parameter is not well-estimated here because it is largely constrained by error responses (Ratcliff & Tuerlinckx, 2002) and the number of error responses was too small.

For the variable-onset task, η was estimated as 0.022. Another fit was performed for the fixed-onset task with η fixed at 0.022 to see how close the fits were for the two tasks if this parameter was the same across tasks. The results are shown in the bottom line of Table S4 and they show that all the parameters except boundary separation are very close to each other. The conclusion is that the two tasks are modeled in the same way and the main difference between them is that boundary separation is larger for the variable-onset task than for the fixed-onset task.

The one-choice task. The one-choice task was essentially a replication of Experiment 1 in the body of the article. The model was fit to the data for individual subjects and the best-fitting parameter values are presented in Table S4. The one-choice model fit the mean RTs and the .1 quantile RTs reasonably well. The parameter values were quite similar to the values for Experiment 1 (Table 1). The ratios of drift rate to SD in drift rate were similar to within a few percent. Nondecision time and boundary separation also matched well.

Figure S5 shows RT histograms for the one-choice task. In order to spread out the tail of the timed-out RTs so they did not bunch up at the 2,000 ms time-out, we added the exponentially distributed stimulus onset time to the 2,000 ms time-out. This is ad hoc, but it would mean that a process that had reached the time out had the same probability of terminating at any point in time given that it had survived to that time. For most of the subjects, the right-most RTs are plausibly long tails. For only three of the subjects is the proportion of responses bunched up just beyond the time-out sufficiently large to suggest that these responses come from failures of detection. The same pattern of results is obtained for the variable onset two-choice task.

Overall, the diffusion model explains the data for the two two-choice tasks in the same way as for the two-choice task in Experiment 1. This generalizes the applicability of the diffusion model to the task with variable stimulus onset time.

Figure Captions

Figure S1. Examples of predicted histograms (frequency polygons, top panel) and hazard functions (bottom panel) for parameter values in the top six lines of Table S1.

Figure S2. Histograms for the data from Experiment 2 in the main article.

Figure S3. Correlations and scatter plots among parameter values for fits of the model with randomly varying starting points to one set of simulated data. The parameter values are similar to those from Experiment 1 and these are shown on the diagonal (along with histograms of the parameter values). The top panel has a low drift rate and the bottom panel has a higher drift rate.

Figure S4. A plot equivalent to that in Figure S3 for fits to simulated data using parameter values similar to those for Experiment 1.

Figure S5. Histograms for the one-choice task from Experiment S1. The ellipses show plausible cases when subjects might have missed the stimulus.

Table S1: One-Choice Diffusion Model Parameter Estimates and Goodness of Fit

Condition	Negative drift rate?	T_{er}	s_t	a	v	η	v/η	a/v	χ^2
Nondeprived	Yes	175.5	103.7	0.092	0.769	0.292	2.63	0.119	20.7
		160.6	104.3	0.126	1.006	0.396	2.53	0.125	22.7
		180.9	104.9	0.083	0.705	0.267	2.63	0.118	21.8
Deprived	Yes	187.2	123.8	0.096	0.500	0.330	1.51	0.192	19.3
		179.2	130.0	0.123	0.654	0.432	1.51	0.188	19.6
		198.7	128.2	0.079	0.412	0.268	1.53	0.192	19.7
Nondeprived	No	167.1	100.7	0.111	0.931	0.372	2.67	0.119	22.1
Deprived	No	180.1	124.2	0.120	0.650	0.472	1.66	0.184	17.6

Note: The three rows for each negative drift rate conditions use different starting values and illustrate the parameter identifiability issue. The two rows for the non-negative drift rate conditions are to show the parameter estimates and goodness of fit for one fit. χ^2, df=14, critical value 26.1. a=boundary setting, T_{er}=nondecision component of response time, η=standard deviation in drift across trials, s_t = range of the distribution of nondecision times, and v is drift rate.

Table S2: Mean Chi-Square Values Across Subjects with Degrees of Freedom

Model: with these variability parameters	Experiment 1 with 3 η's	Experiment 1 with one η	Experiment 2 test bout by test bout	Experiment 2 with grouped data
η , s_a, s_t	48.8 (49)	52.1 (51)	13.9 (13)	18.1 (13)
η , s_t	48.0 (50)	53.4 (52)	14.1 (14)	20.0 (14)
s_t	66.3 (51)	66.3 (53)	25.7 (15)	87.8 (15)
η	55.6 (51)	59.2 (53)	16.8 (15)	38.7 (15)
s_a, s_t				92.2 (14)
η , s_a				36.2 (14)
s_a				129.2 (15)

Note, η=standard deviation in drift across trials, s_t = range of the distribution of nondecision times, s_a is the range of the distribution of boundary position.

Table S3: Data and One-Choice Diffusion Model Predictions from Experiment S1

Task	Proportion of white or black pixels	Data			Theory			Data
		Accuracy	Mean RT	0.1 quantile RT	Accuracy	Mean RT	0.1 quantile RT	Proportion time-out
One-choice variable onset	.60		516	364		524	355	.004
	.56		802	445		772	440	.072
Two-choice variable onset	.60	.961	596	436	.999	585	438	.002
	.56	.978	842	516	.967	851	523	.068
Two-choice fixed onset	.60	.971	541	418	.996	527	421	0
	.56	.957	627	441	.969	611	444	0

Table S4: Diffusion Model Parameters

Task	T_{er1}	T_{er2}	s_t	a	z	s_z	v_1	v_2	η_1	η_2
One-choice	254		187	.136			.648	.336	.214	.164
Two-choice fixed onset	360	364	105	.160	.081	.027	.519	.356	.150	
Two-choice variable onset	335	360	137	.192	.097	.035	.384	.185	.022	
Two-choice fixed onset (η=.022)	364	370	92	.125	.063	.004	.379	.235	.022	

Note. a=boundary setting, T_{er}=nondecision component of response time, η=standard deviation in drift across trials, s_t = range of the distribution of nondecision times, and v is drift rate. Subscript 1 is for the easy stimulus condition and subscript 2 is for the difficult stimulus condition.

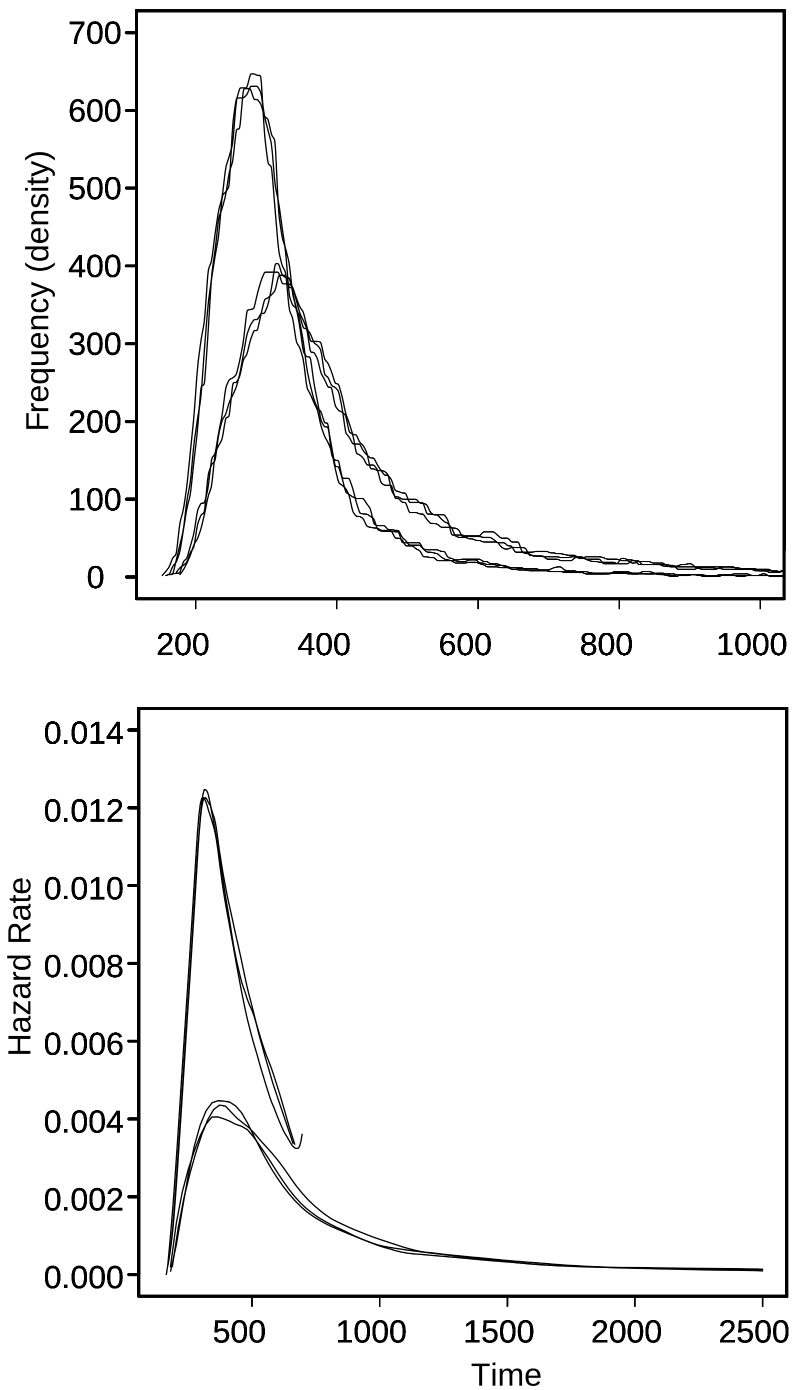

Figure S1. Examples of predicted histograms (frequency polygons, top panel) and hazard functions (bottom panel) for parameter values in the top six lines of Table S1.

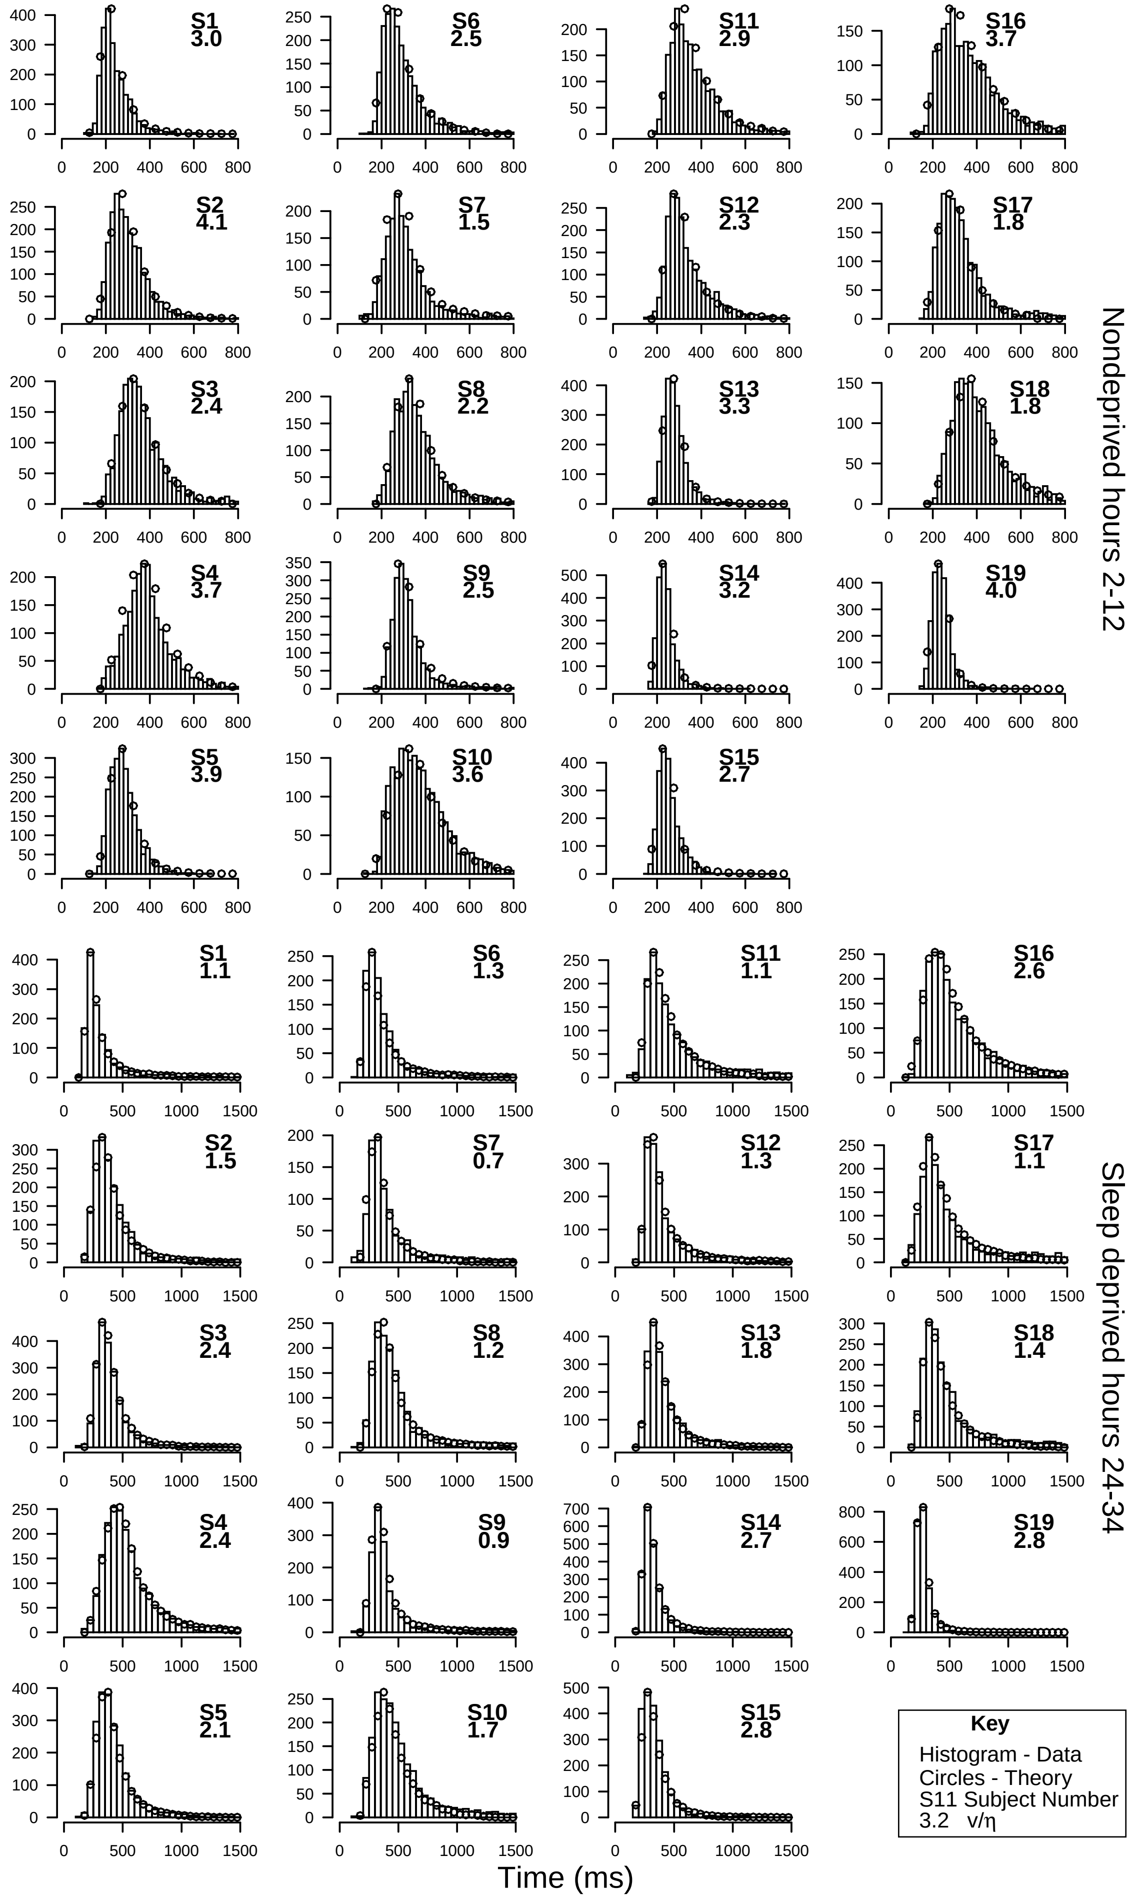

Figure S2. Histograms for the data from Experiment 2 in the main article.

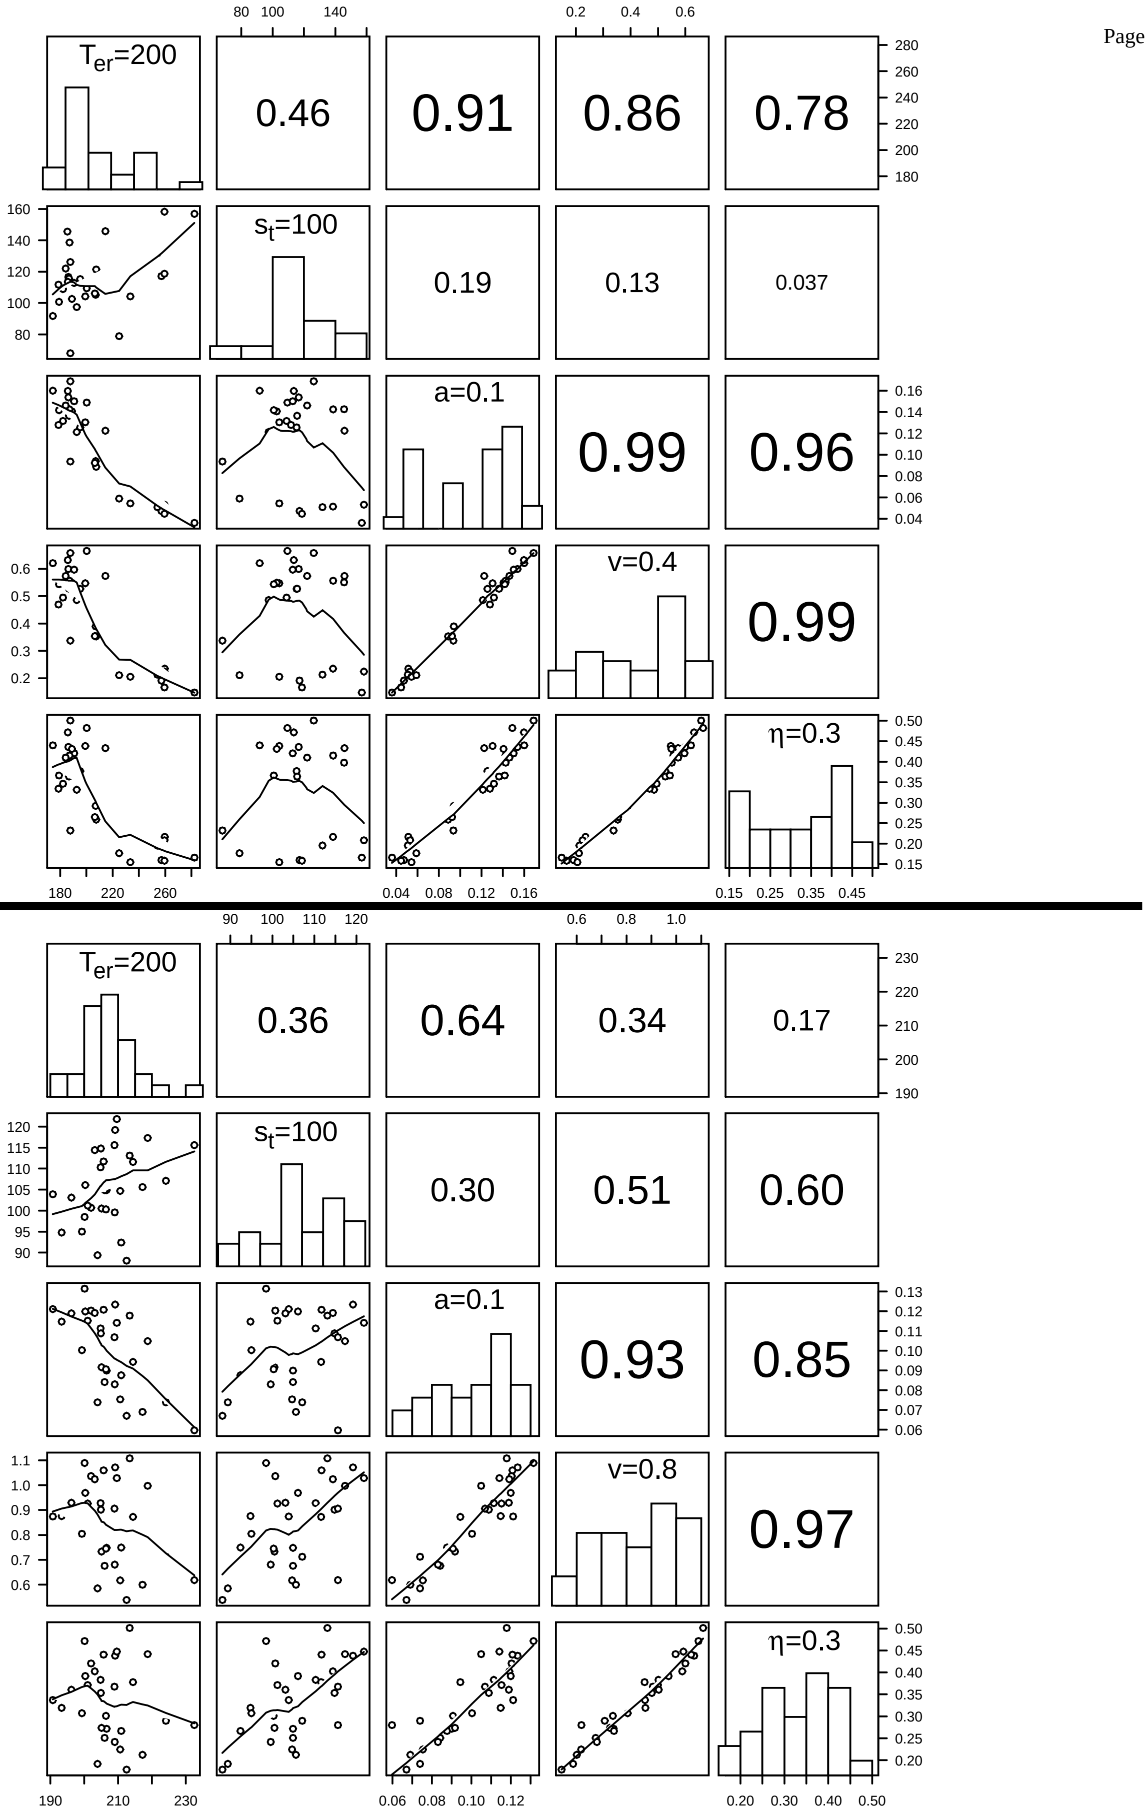

Figure S3. Correlations and scatter plots among parameter values for fits of the model with randomly varying starting points to one set of simulated data. The parameter values are similar to those from Experiment 1 and these are shown on the diagonal (along with histograms of the parameter values). The top panel has a low drift rate and the bottom panel has a higher drift rate

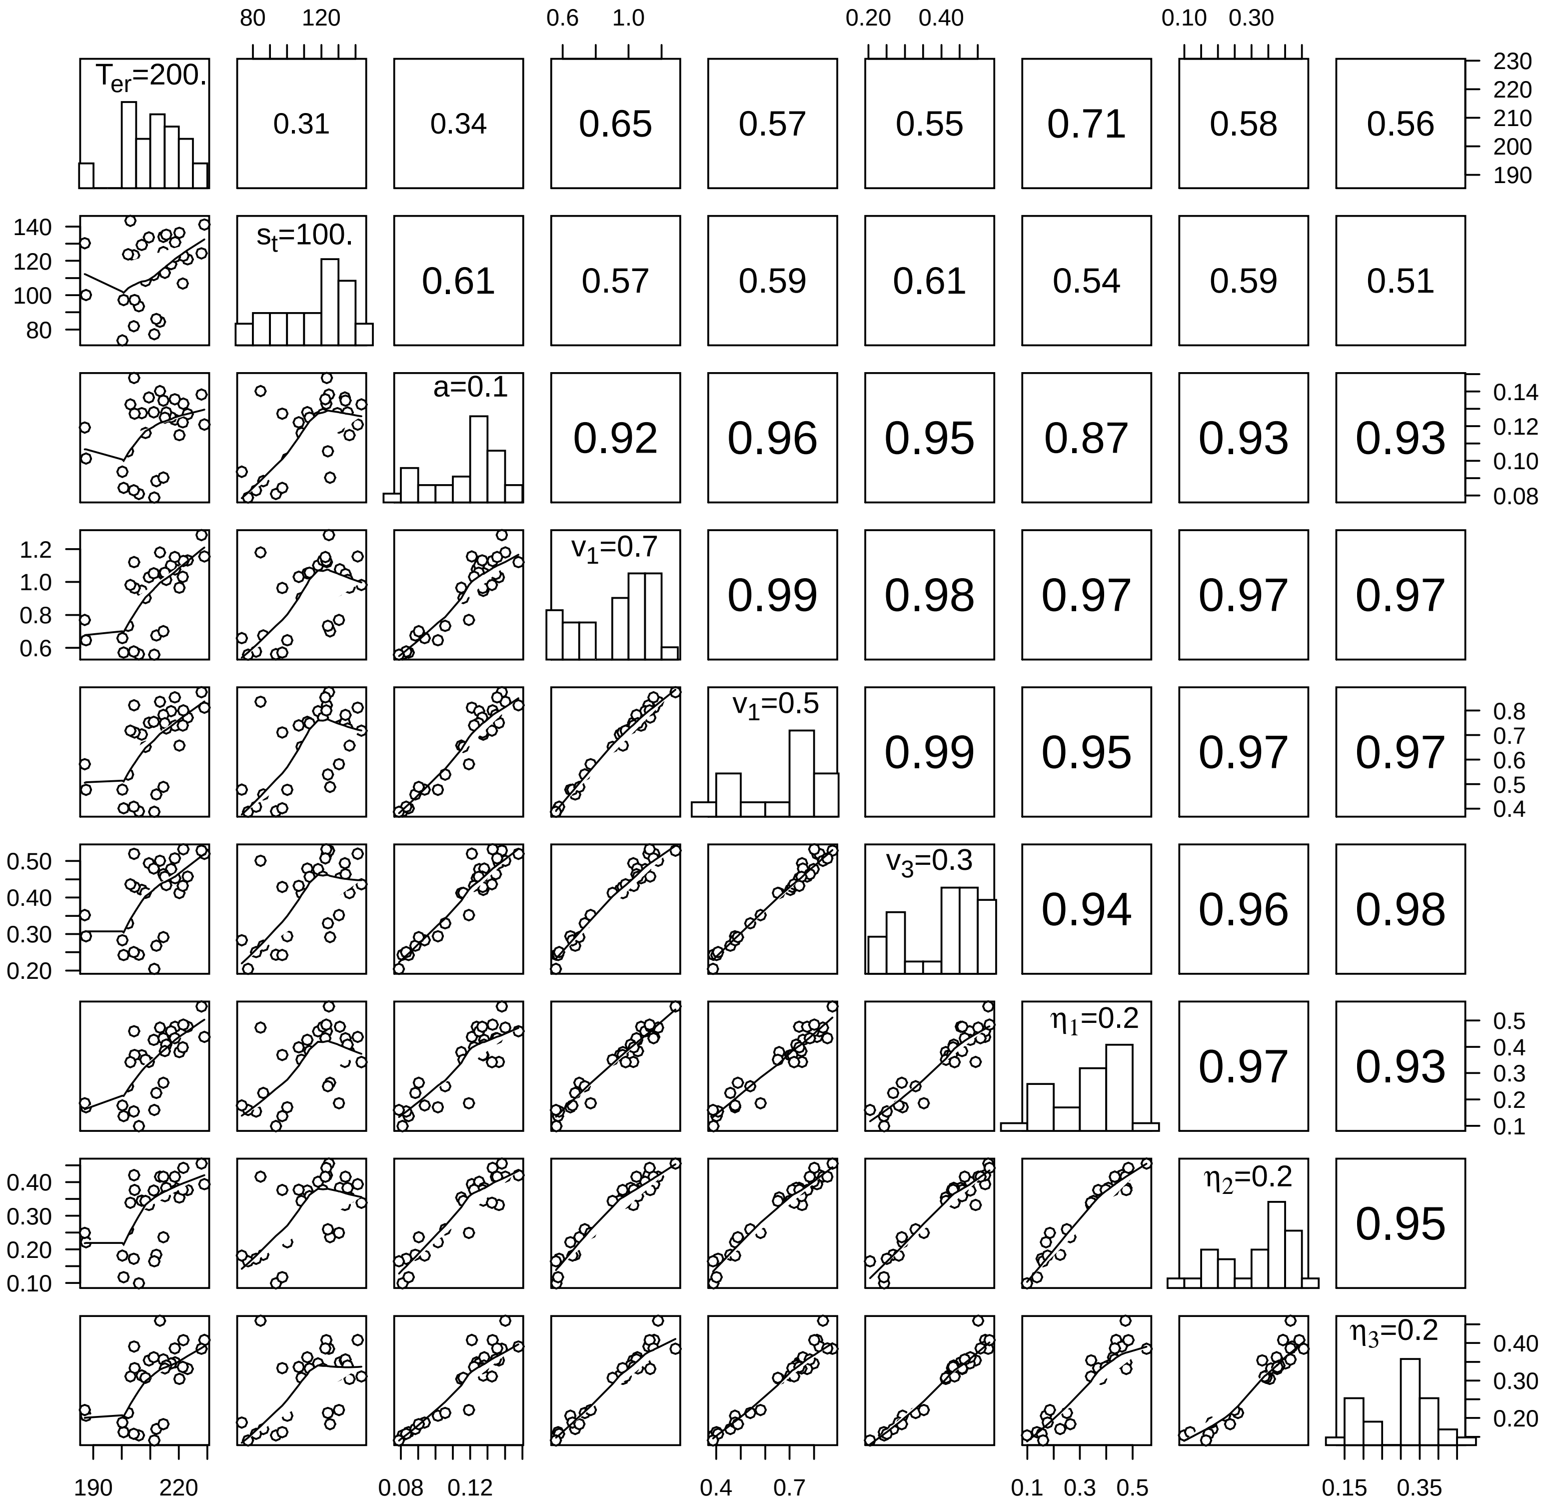

Figure S4. A plot equivalent to that in Figure S3 for fits to simulated data using parameter values similar to those for Experiment 1.

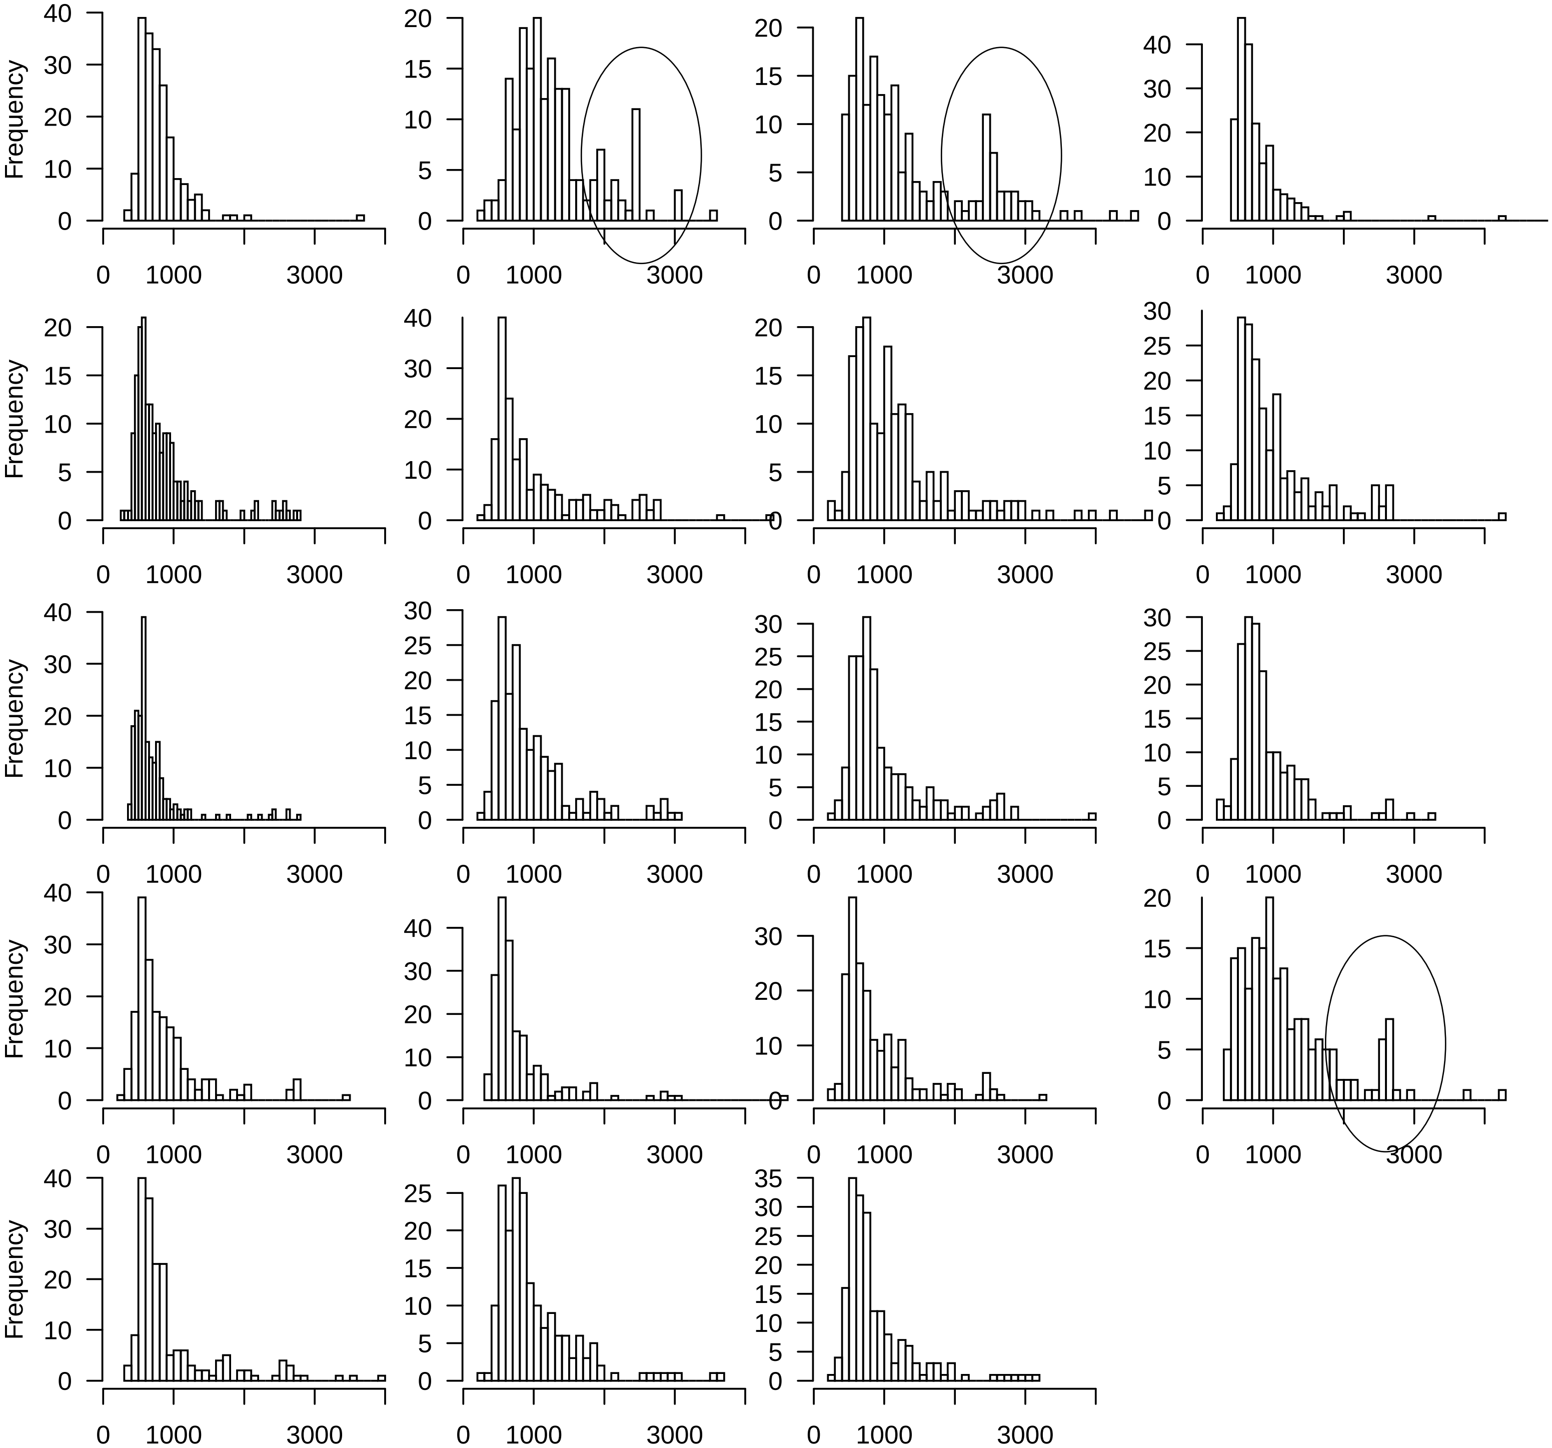

Figure S5. Histograms for the one-choice task from Experiment S1.
The ellipses show plausible cases when subjects might have missed the stimulus.