Supplemental materials for *Base rate neglect and conservatism in probabilistic reasoning: Insights from eliciting full distributions*

Piers Douglas Lionel Howe[1], Andrew Perfors[1], Bradley Walker[2],
Yoshihisa Kashima[1], and Nicolas Fay[2]
[1]School of Psychological Sciences, University of Melbourne
[2]School of Psychological Science, University of Western Australia

Experiment 1

Main analyses

Figure 1

Experiment 1: Relationship between β and γ of individuals after both five chips and an unlimited number of chips

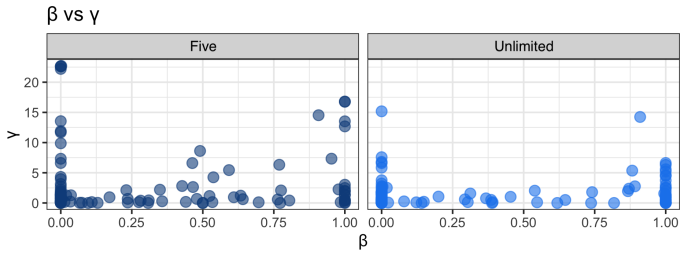

Note. Relationship between fitted β and γ values in the MAIN condition of Experiment 1. Each dot indicates one participant. The plots show that after both five and unlimited chips, there is no systematic relationship between β and γ.

Figure 2

Experiment 1: Goodness of model fit after five chips and after an unlimited number of chips in the MAIN *condition*

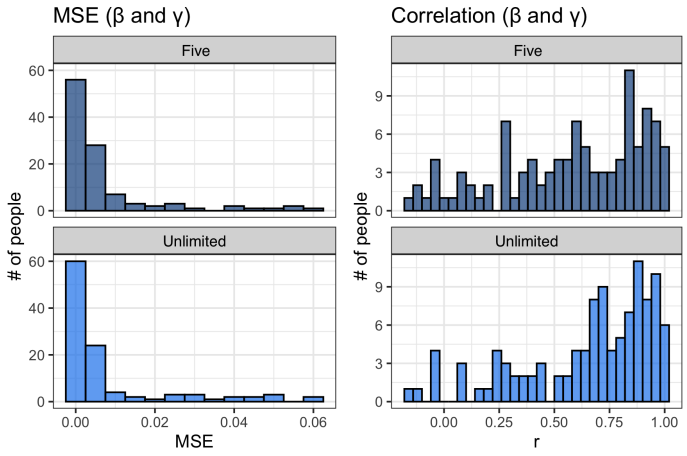

Note. We quantify how well our model fits individuals in two ways. The panel on the left shows the distribution of mean squared error (MSE) between the fitted posteriors (i.e. the posteriors generated by the model) and the participants' reported posteriors. It is evident that most fits were good, with MSE values below 0.01. The panel on the right shows the correlation between the fitted posteriors and the participants' reported posteriors. Again, most correlations (r) were strong, indicating that most participants were fit well.

Good participants

Below we report all analyses on the subset of participants (76 of the original 107) who were fit well by the model, i.e., who had an MSE less than 0.01. We'll call these the GOOD ones.

Figure 3

Experiment 1 GOOD *participants: Aggregate best-fit estimates*

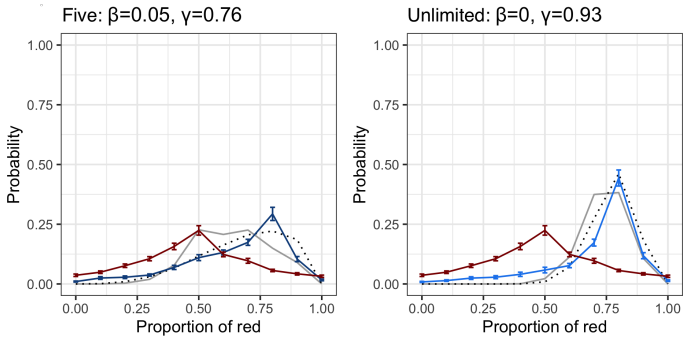

Note. Reported distributions for prior beliefs (red line) as well as after seeing five (dark blue line, left panel) and unlimited (light blue line, right panel) chips. In both panels, the grey line indicates the optimal Bayesian prediction given the prior, while the black dotted line indicates the prediction of the line of best fit based on the inferred parameters β and γ. In both panels, as in the full dataset, β is around zero, indicating that the aggregate posterior was best fit assuming that participants disregarded their reported prior (i.e., showed almost complete base rate neglect). The value for γ indicates some conservatism on average in both conditions, somewhat less than was observed in the full dataset.

Figure 4

Experiment 1 GOOD *participants: Relationship between* β *and* γ

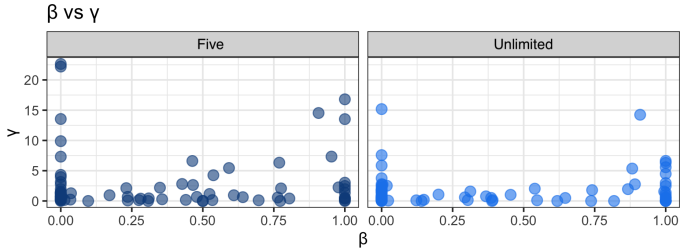

Note. Relationship between obtained β and γ values. Each dot indicates one participant. The plots show that after both five and unlimited chips, there is no systematic relationship between β and γ.

Figure 5

Experiment 1 GOOD *participants: Distribution of fitted* β *and* γ *values across individuals*

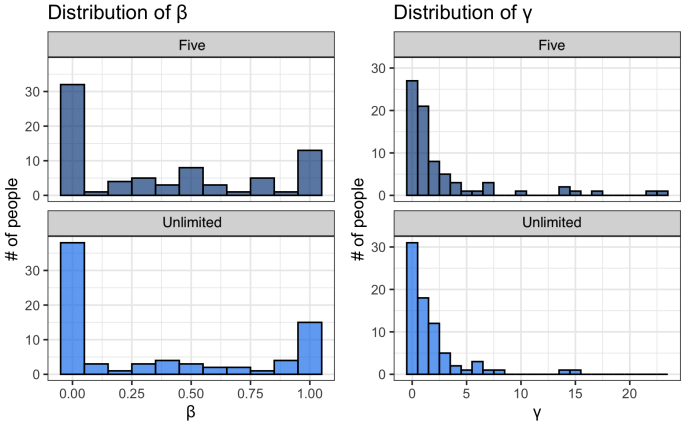

Note. Histograms showing the distribution of best-fit β and γ values across individuals. The panels on the left show that the majority of people showed a moderate or large amount of base rate neglect; their inferences were best described with β values less than one and often close to zero. 43.4% of people after seeing five chips and 50% after seeing an unlimited number of additional chips had $\beta < 0.1$, and 18.4% after seeing five chips and 21.1% after seeing an unlimited number had $\beta > 0.9$. The value for γ indicates a range of conservatism, with 52.6% after seeing five chips and 57.9% after seeing an unlimited number with $\gamma < 1$.

Figure 6

Experiment 1 GOOD *participants: Distribution of fitted* β *and* γ *holding the other constant*

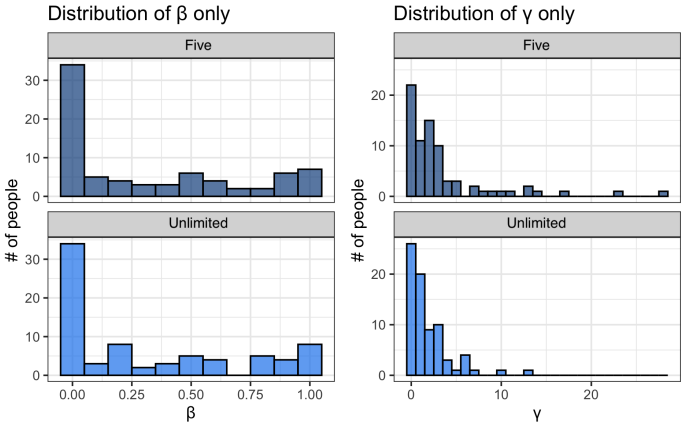

Note. Histograms showing the distribution of best-fit β and γ values across individuals, holding the other constant ($\gamma = 1$ while fitting β, and $\beta = 1$ while fitting γ). As in the full dataset, most people's inferences were best described with β values less than one and often close to zero (34.6% of people after five chips and 31.8% after unlimited chips had $\beta < 0.1$, and 8.4% after both five chips and unlimited chips had $\beta > 0.9$). The value for γ indicates somewhat less conservatism in both conditions, with 26.2% of people after five chips and 36.4% after unlimited chaps having a $\gamma < 1$.

Experiment 2

Main analyses

Figure 7

Experiment 2: Relationship between β and γ

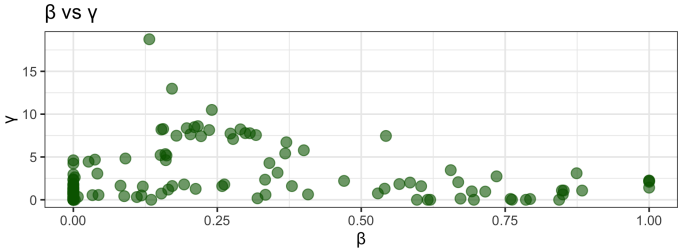

Note. Relationship between β and γ values for the PEAKED condition. Each dot indicates one participant. There is no systematic relationship between β and γ, but the distribution of the two parameters is strikingly different than in the other conditions. There were very few participants with extremely high γ values (above 15 or so), unlike in the other experiments. In those experiments, a peaked posterior at 80% could only be accounted for with this very high γ value, since β changed the shape of the distribution. In Experiment 2, however, a peaked posterior at 80% could also be accounted for with a somewhat more moderate γ value (above 1 but less than 15) in conjunction with a non-zero β. This is one indication of the non-identifiability problem discussed in the main paper. (The UNIFORM condition is not shown because β is not inferred there.)

Figure 8

Experiment 2: Goodness of model fit

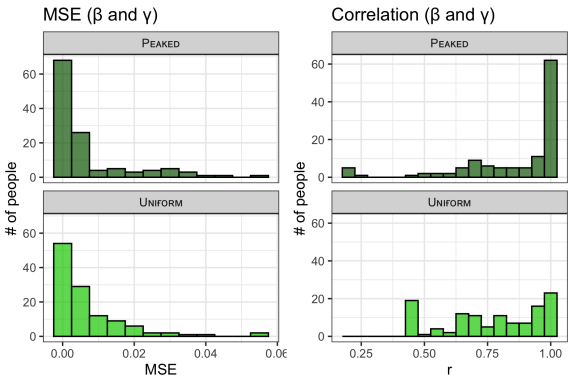

Note. The panel on the left shows the distribution of mean squared error (MSE) between the fitted posteriors and the reported posteriors. The panel on the right shows the correlation between the fitted and reported posteriors. Fits in both conditions were good but noticeably better fit in the PEAKED condition, presumably because it had two free parameters instead of one.

Good participants

Below we report all analyses on the subset of participants (189 of the original 239) who were fit well by the model, i.e., who had an MSE less than 0.01. As before, we'll call these the GOOD ones.

Figure 9

Experiment 2 GOOD *participants: Aggregate best-fit estimates*

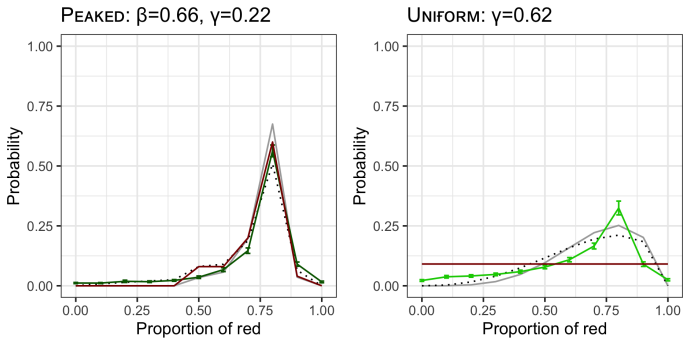

Note. Reported distributions in the PEAKED (dark green line, left panel) and UNIFORM (light green line, right panel) conditions. The grey line indicates the optimal Bayesian prediction given the prior, while the black dotted line indicates the prediction of the line of best fit based on the inferred parameters β and γ. In the PEAKED condition, β indicates a moderate degree of base rate neglect: higher than in Experiment 1 but still present. The value for γ again indicates moderate conservatism, although given the identifiability issues these parameters should be interpreted with caution.

Figure 10

Experiment 2 GOOD *participants: Relationship between* β *and* γ

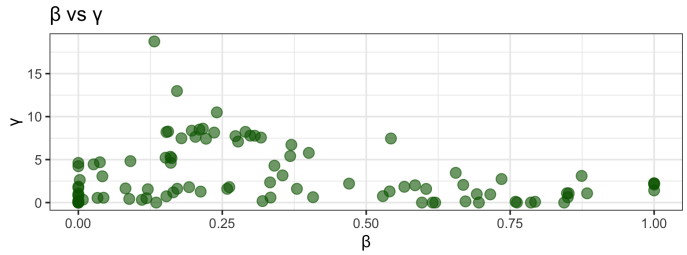

Note. Relationship between obtained β and γ values. Each dot indicates one participant. As in the full dataset, there was no systematic relationship between β and γ.

Figure 11

Experiment 2 GOOD *participants: Distribution of fitted* β *and* γ *values across individuals*

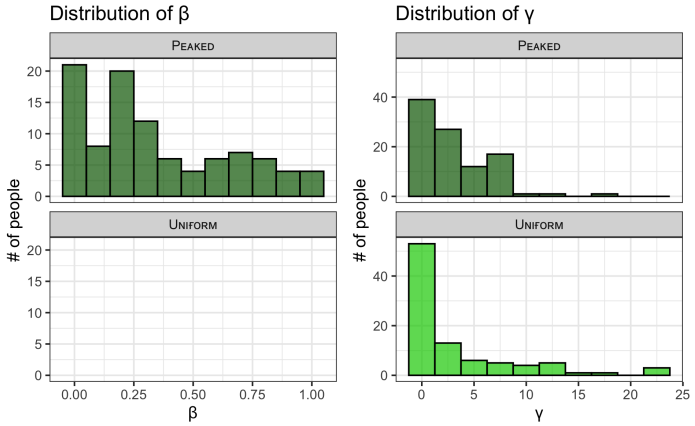

Note. Histograms showing the distribution of best-fit β and γ values across individuals in the PEAKED and UNIFORM conditions. The panels on the left show that the majority of people showed some amount of base rate neglect, but less *extreme* base rate neglect in this sample than in the full dataset (24.5% of people had $\beta < 0.1$, compared to 38.8% in the full dataset); however, only a small minority showed none at all (4.1% had $\beta > 0.9$, compared to 3.3% in the full dataset). The value for γ indicates a range of conservatism in both conditions (though less in PEAKED), with 34.7% in PEAKED and 54.9% in UNIFORM with $\gamma < 1$. This probably results, at least in part, from the unidentifiability issues discussed elsewhere.

Figure 12

Experiment 2 GOOD *participants: Distribution of fitted* β *and* γ *holding the other constant*

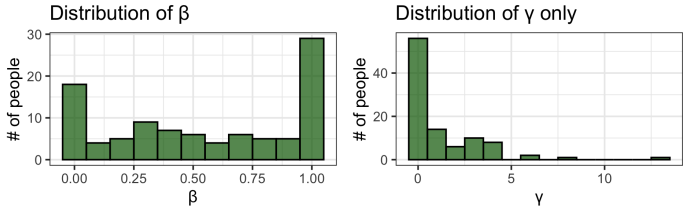

Note. Histograms showing the distribution of best-fit β and γ values across individuals in the PEAKED condition, holding the other parameter constant (inferring β when $\gamma = 1$ and γ when $\beta = 1$). As with the full dataset, these parameter distributions look different than the joint ones, but even more so. Now more people showed no base rate neglect than complete base rate neglect (19.4% of people had $\beta < 0.1$, and 31.6% had $\beta > 0.9$). The value for γ indicates that many more were conservative (68.4% had $\gamma < 1$) when β was constrained to 1. This is more evidence of the unidentifiability issue, as people with moderate γ and low β in the joint distribution were fit very differently when done separately.

Experiment 3

Main analyses

Figure 13

Experiment 3: Relationship between β and γ

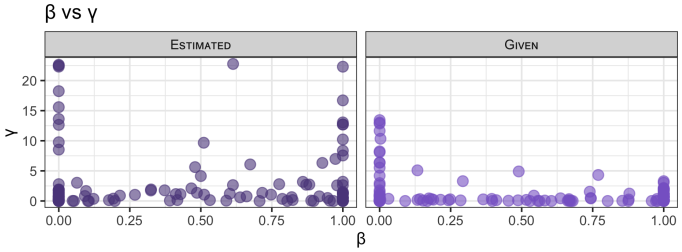

Note. Relationship between obtained β and γ values. Each dot indicates one participant. The plots show that in both ESTIMATED and GIVEN conditions, there is no systematic relationship between β and γ.

Figure 14

Experiment 3: Goodness of fit

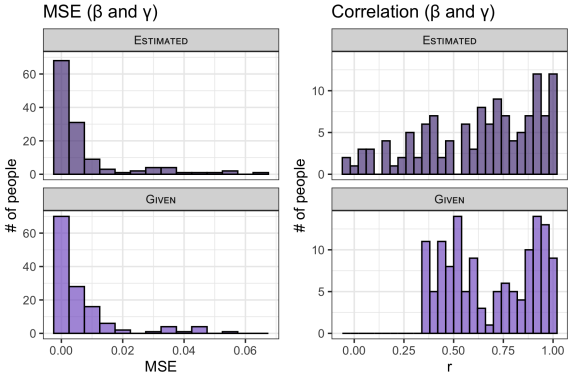

Note. We quantify how well our model fits individuals in two ways. The panel on the left shows the distribution of mean squared error (MSE) between the fitted posteriors and the reported posteriors. It is evident that most fits were good, with MSE values below 0.01. The panel on the right shows the correlation between the fitted posteriors and the reported posteriors. Again, most correlations (r) were strong, indicating that most participants were fit well.

Good participants

Below we report all analyses on the subset of participants (213 of the original 261) who were fit well by the model, i.e., who had an MSE less than 0.01. As before, we'll call these the GOOD ones.

Figure 15

Experiment 3 GOOD *participants: Aggregate best-fit estimates*

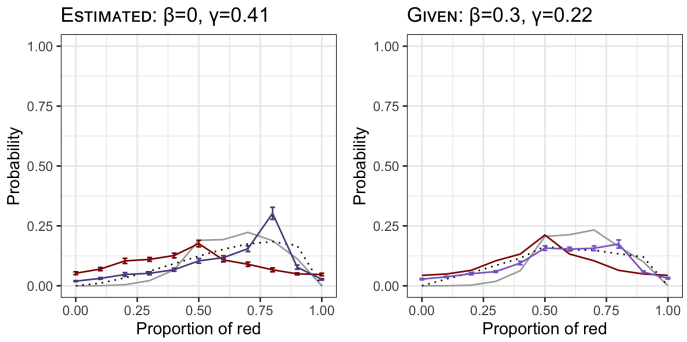

Note. Reported distributions for prior beliefs (red line) and posteriors in the ESTIMATED (dark purple line, left panel) and GIVEN (light purple line, right panel) conditions. In both panels, the grey line indicates the optimal Bayesian prediction given the prior, while the black dotted line indicates the prediction of the line of best fit based on the inferred parameters β and γ. As in the full dataset, β is around zero in the ESTIMATED condition and slightly higher in the GIVEN condition. The value for γ indicates a moderate degree of conservatism on average in both conditions.

Figure 16

Experiment 3 GOOD *participants: Relationship between* β *and* γ

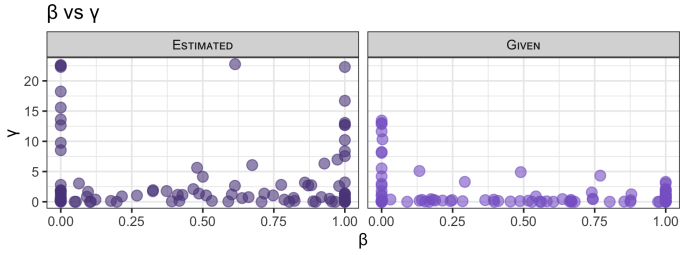

Note. Relationship between obtained β and γ values. Each dot indicates one participant. The plots show that in both ESTIMATED and GIVEN conditions, there is no systematic relationship between β and γ.

Figure 17

Experiment 3 GOOD *participants: Distribution of fitted* β *and* γ *values across individuals*

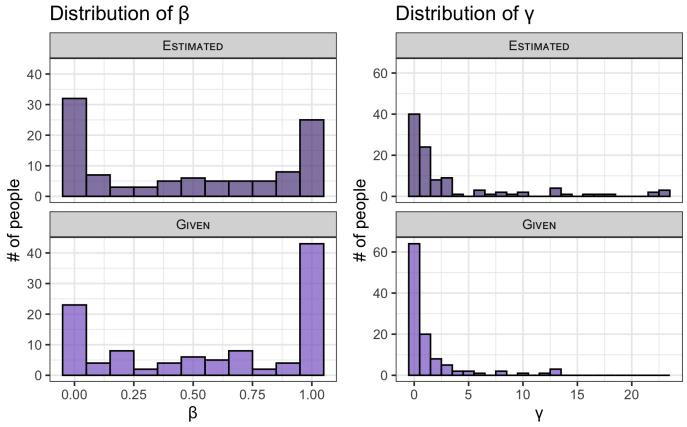

Note. Histograms showing the distribution of best-fit β and γ values across individuals. As in the full dataset, many people showed a moderate or large amount of base rate neglect in the ESTIMATED condition, in the GIVEN condition many people showed no neglect at all. 34.6% of people in ESTIMATED and 22% in GIVEN had $\beta < 0.1$, and 26.9% in ESTIMATED and 39.4% in GIVEN had $\beta > 0.9$. The value for γ indicates a range of conservatism, with 48.1% in ESTIMATED and 68.8% in GIVEN with $\gamma < 1$.

Figure 18

Experiment 3 GOOD *participants: Distribution of fitted* β *and* γ *holding the other constant*

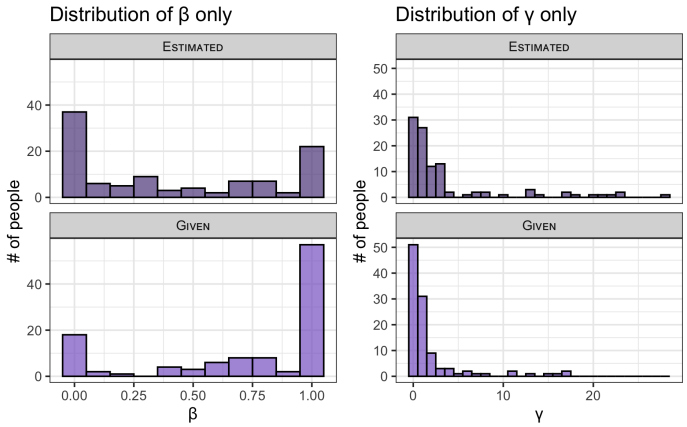

Note. Histograms showing the distribution of best-fit β and γ values across individuals, holding the other constant ($\gamma = 1$ while fitting β, and $\beta = 1$ while fitting γ). The patterns were very similar as found in the full dataset and when the parameters were inferred jointly. 38.5% of people in ESTIMATED and 18.3% in GIVEN had $\beta < 0.1$, and 22.1% in ESTIMATED and 52.3% in GIVEN had $\beta > 0.9$. The value for γ indicates a range of conservatism, with 41.3% in ESTIMATED and 64.2% in GIVEN with $\gamma < 1$.

Figure 19

Experiment 1 GOOD *participants: Each person after five chips, fitting* β *and* γ

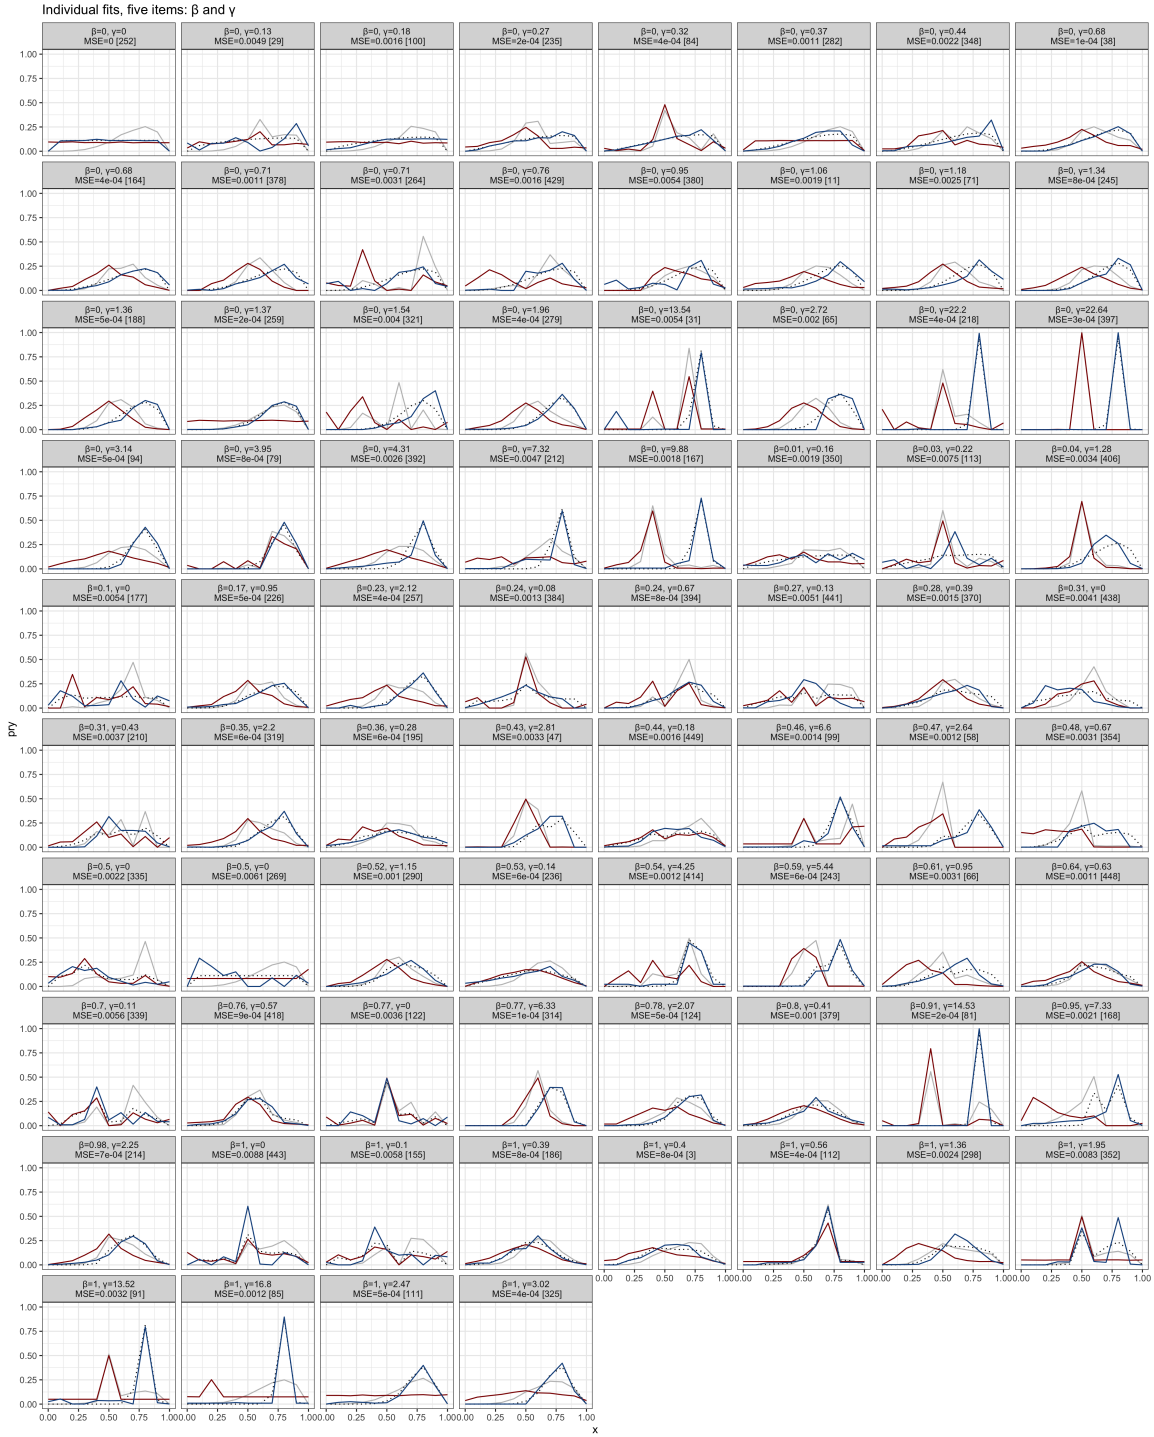

Note. Each good individual participant after receiving five chips. The red line is the reported prior, the dark blue line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 20

Experiment 1 GOOD *participants: Each person after unlimited chips, fitting* β *and* γ

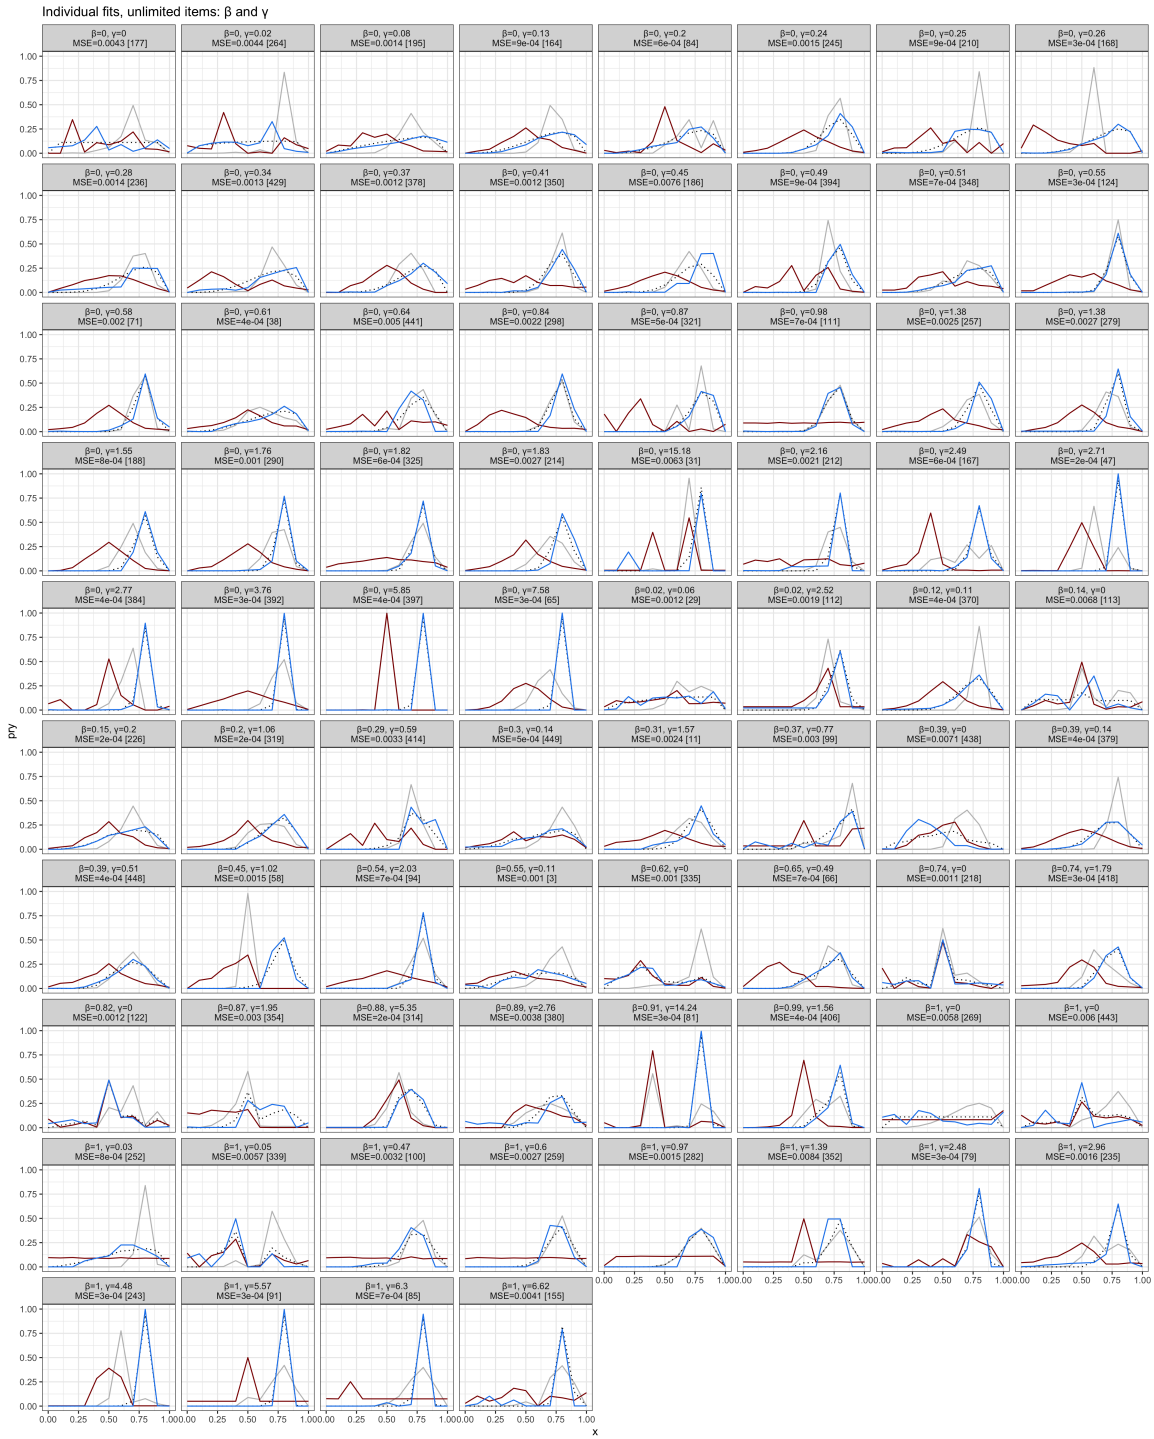

Note. Each good individual participant after receiving unlimited chips. The red line is the reported prior, the light blue line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 21

Experiment 1 BAD *participants: Each person after five chips, fitting* β *and* γ

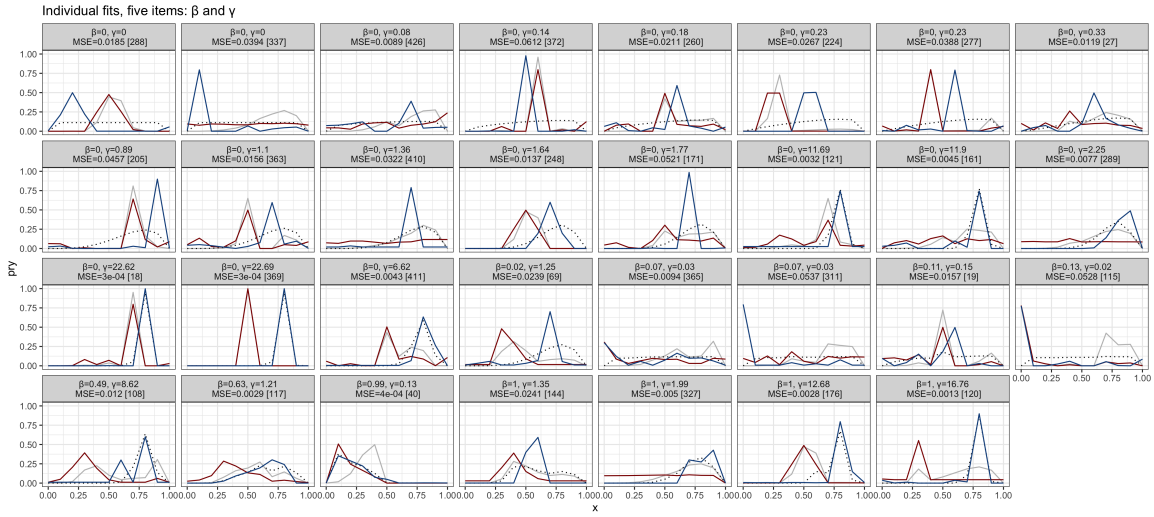

Note. Each bad individual participant after receiving five chips. The red line is the reported prior, the dark blue line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 22

Experiment 1 BAD *participants: Each person after unlimited chips, fitting* β *and* γ

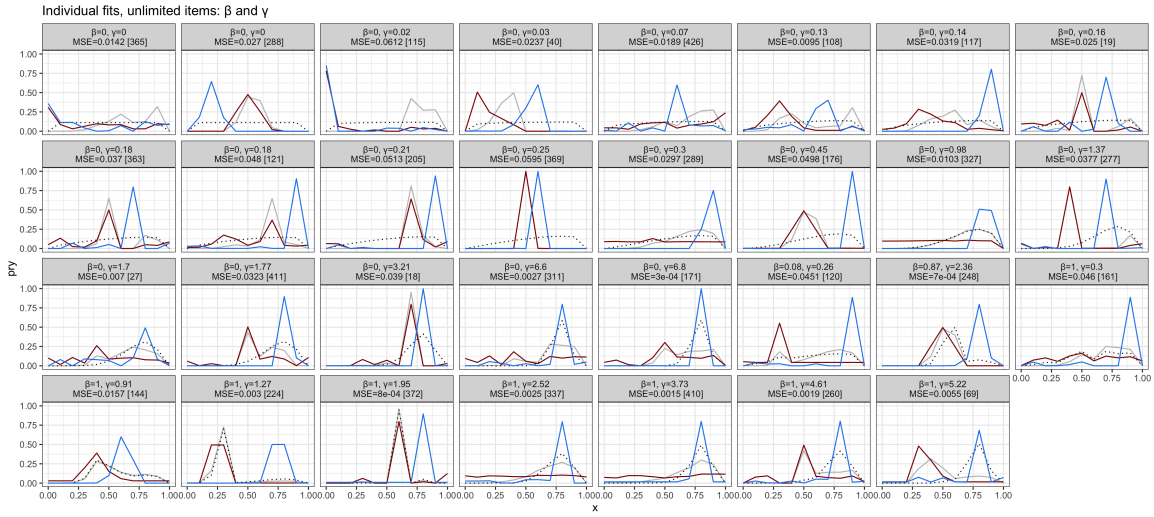

Note. Each bad individual participant after receiving unlimited chips. The red line is the reported prior, the light blue line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 23

Experiment 2 GOOD *participants: Each person in the* PEAKED *condition, fitting* β *and* γ

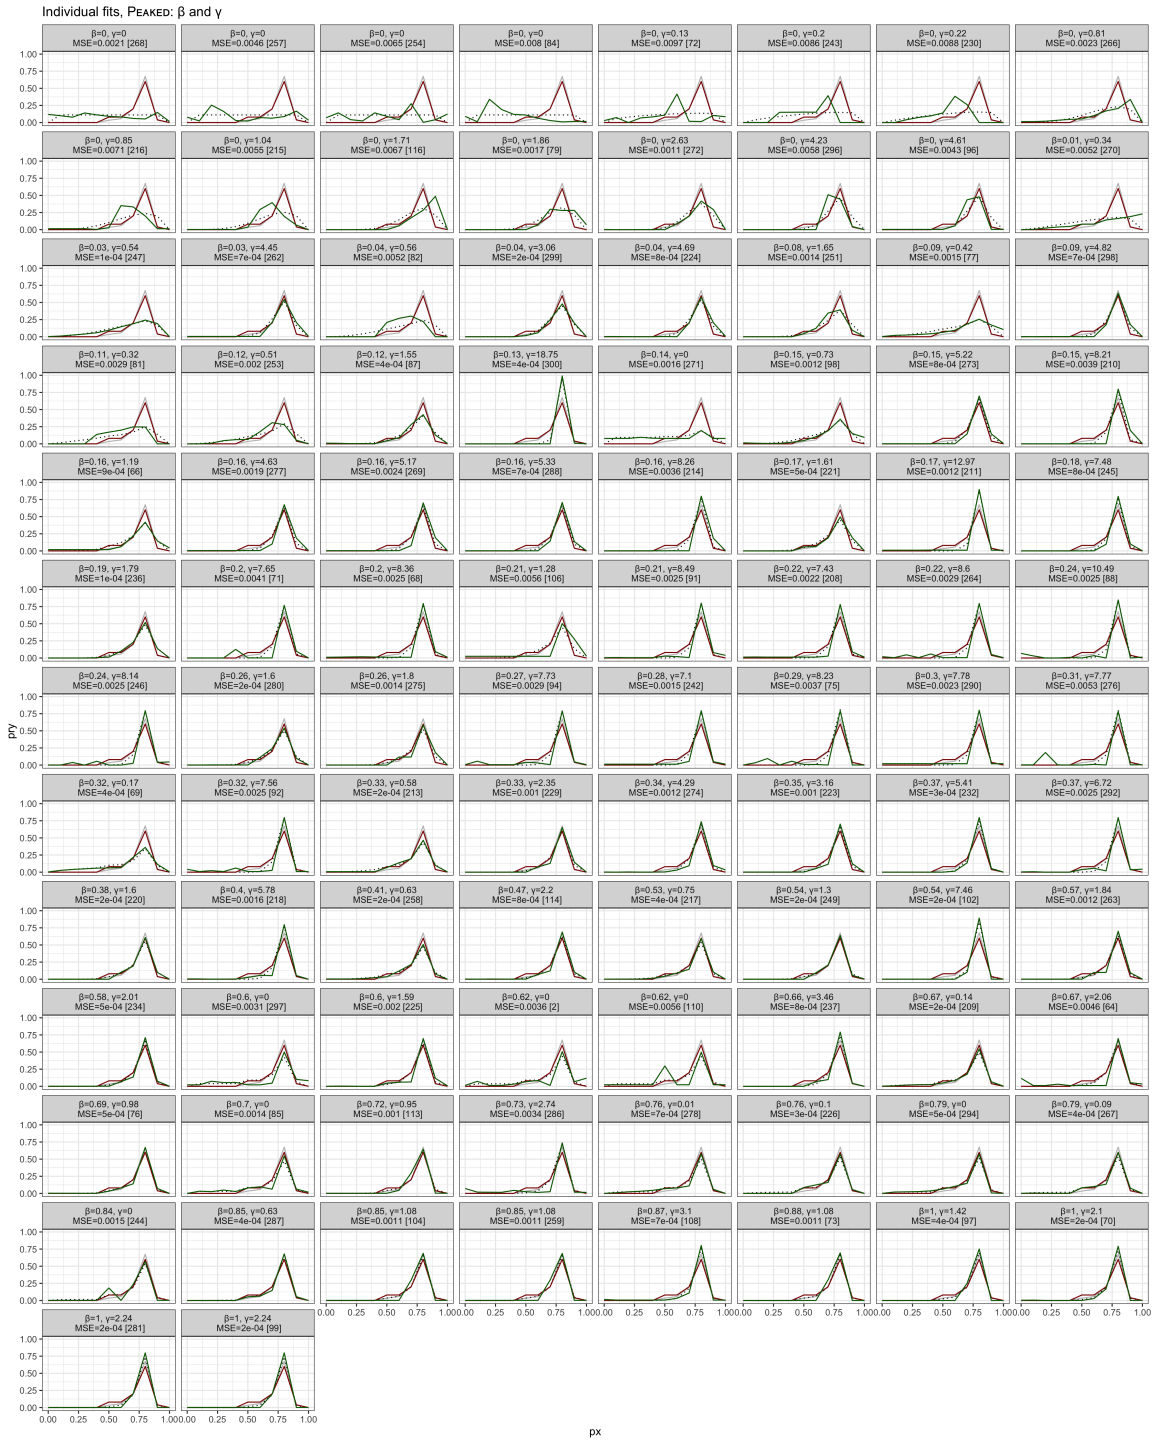

Note. Each good individual participant in the PEAKED condition. The red line is the reported prior, the dark green line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 24

Experiment 2 GOOD *participants: Each person in the* UNIFORM *condition, fitting only* γ

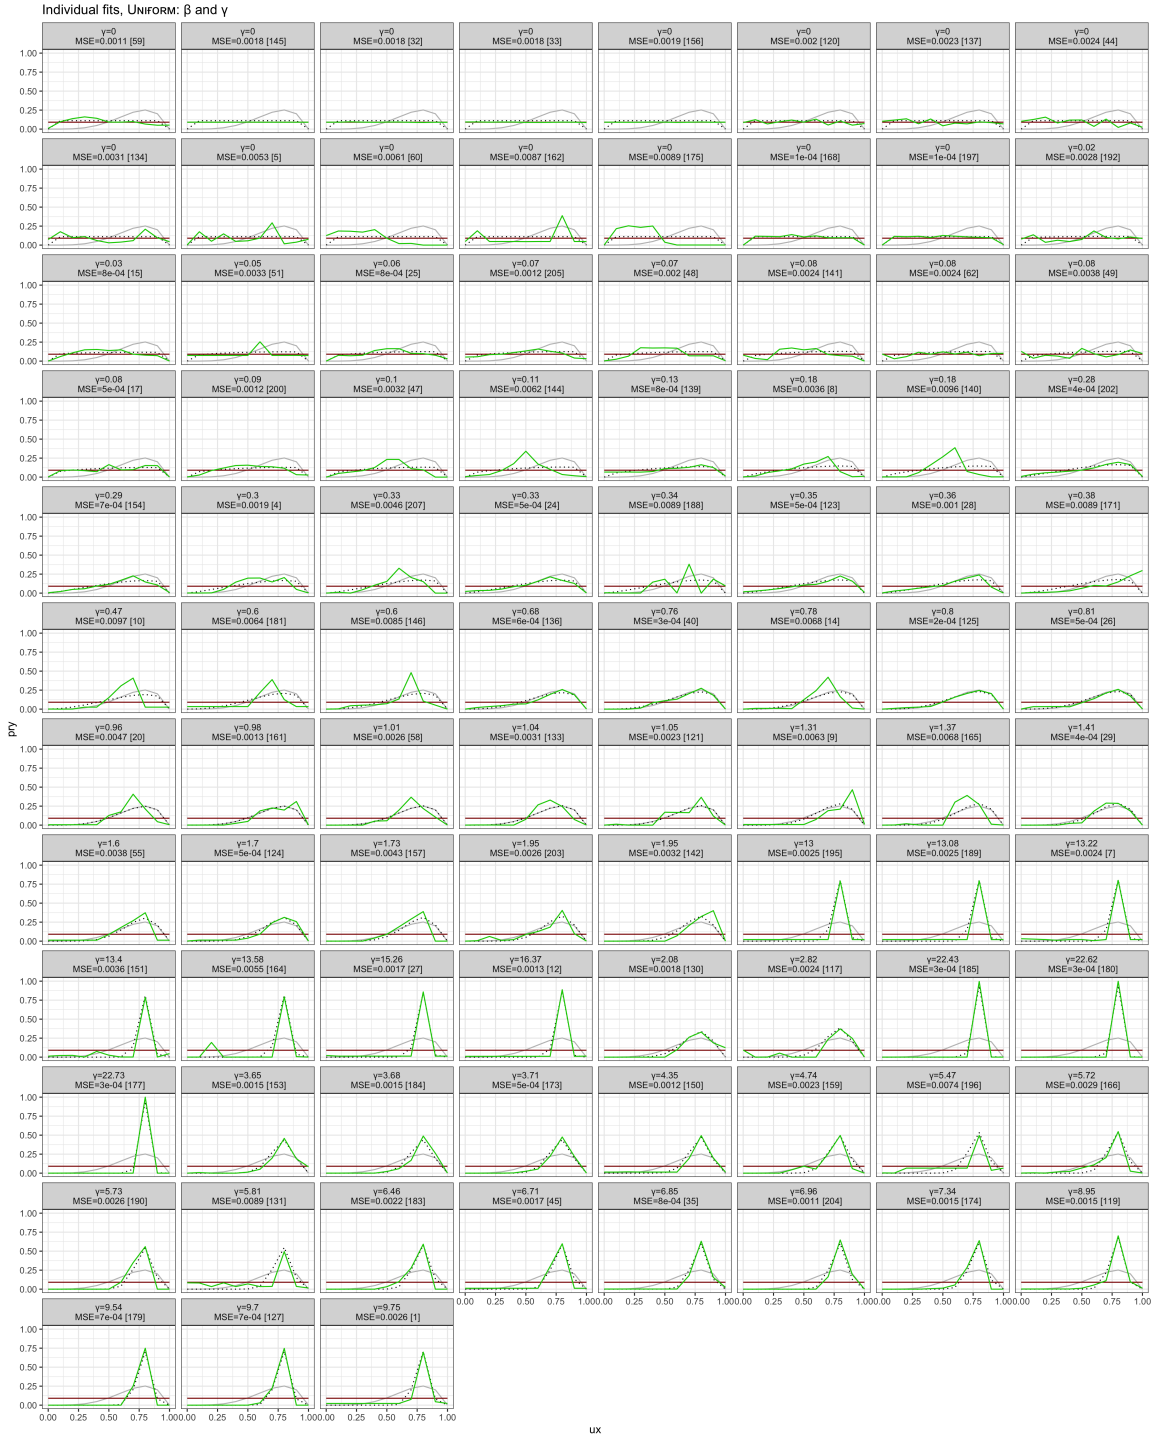

Note. Each good individual participant in the UNIFORM condition. The red line is the reported prior, the light green line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 25

Experiment 2 BAD *participants: Each person in the* PEAKED *condition, fitting* β *and* γ

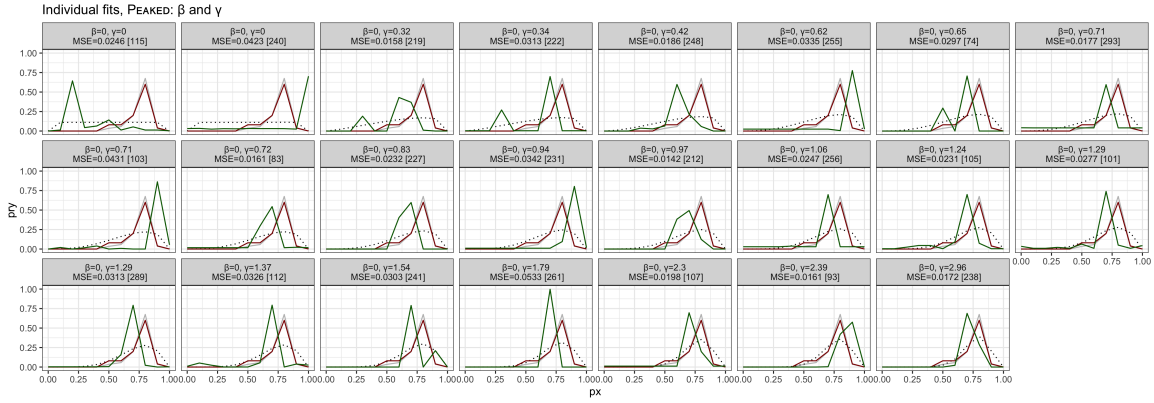

Note. Each bad individual participant in the PEAKED condition. The red line is the reported prior, the dark green line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 26

Experiment 2 BAD *participants: Each person in the* UNIFORM *condition, fitting only* γ

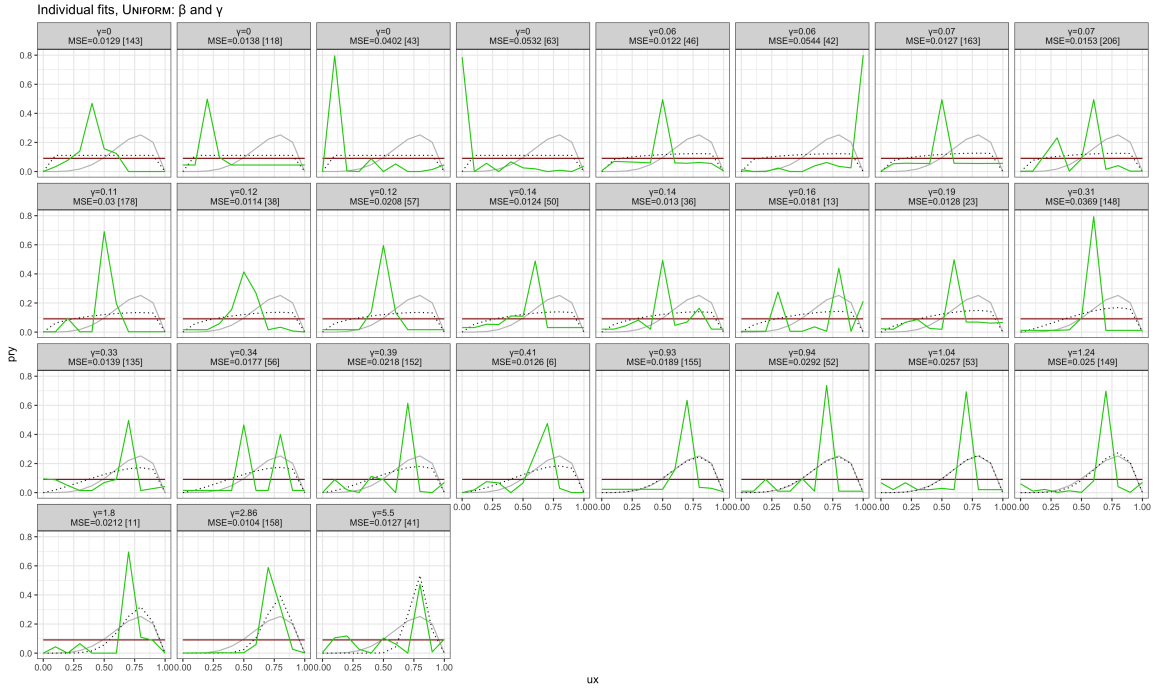

Note. Each bad individual participant in the UNIFORM condition. The red line is the reported prior, the light green line is reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 27

Experiment 3 GOOD *participants: Each person in the* ESTIMATED *condition, fitting* β *and* γ

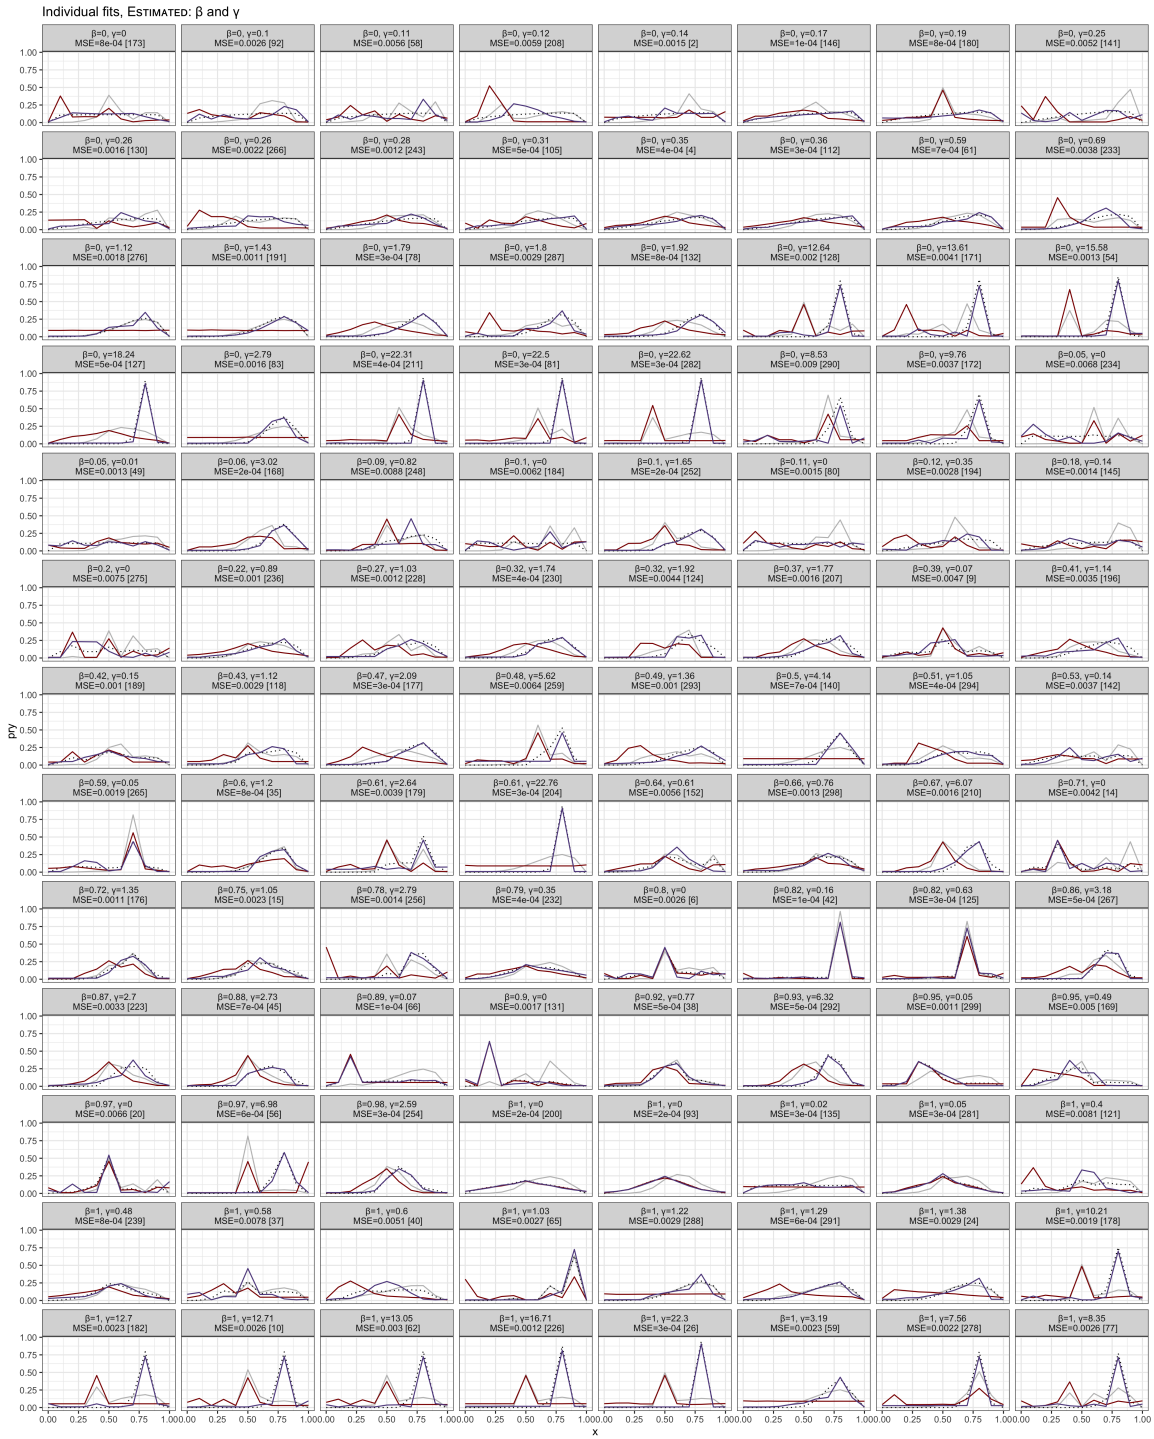

Note. Each good individual participant in the ESTIMATED condition. The red line is the reported prior, the dark purple line is the reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 28

Experiment 3 GOOD *participants: Each person in the* GIVEN *condition, fitting* β *and* γ

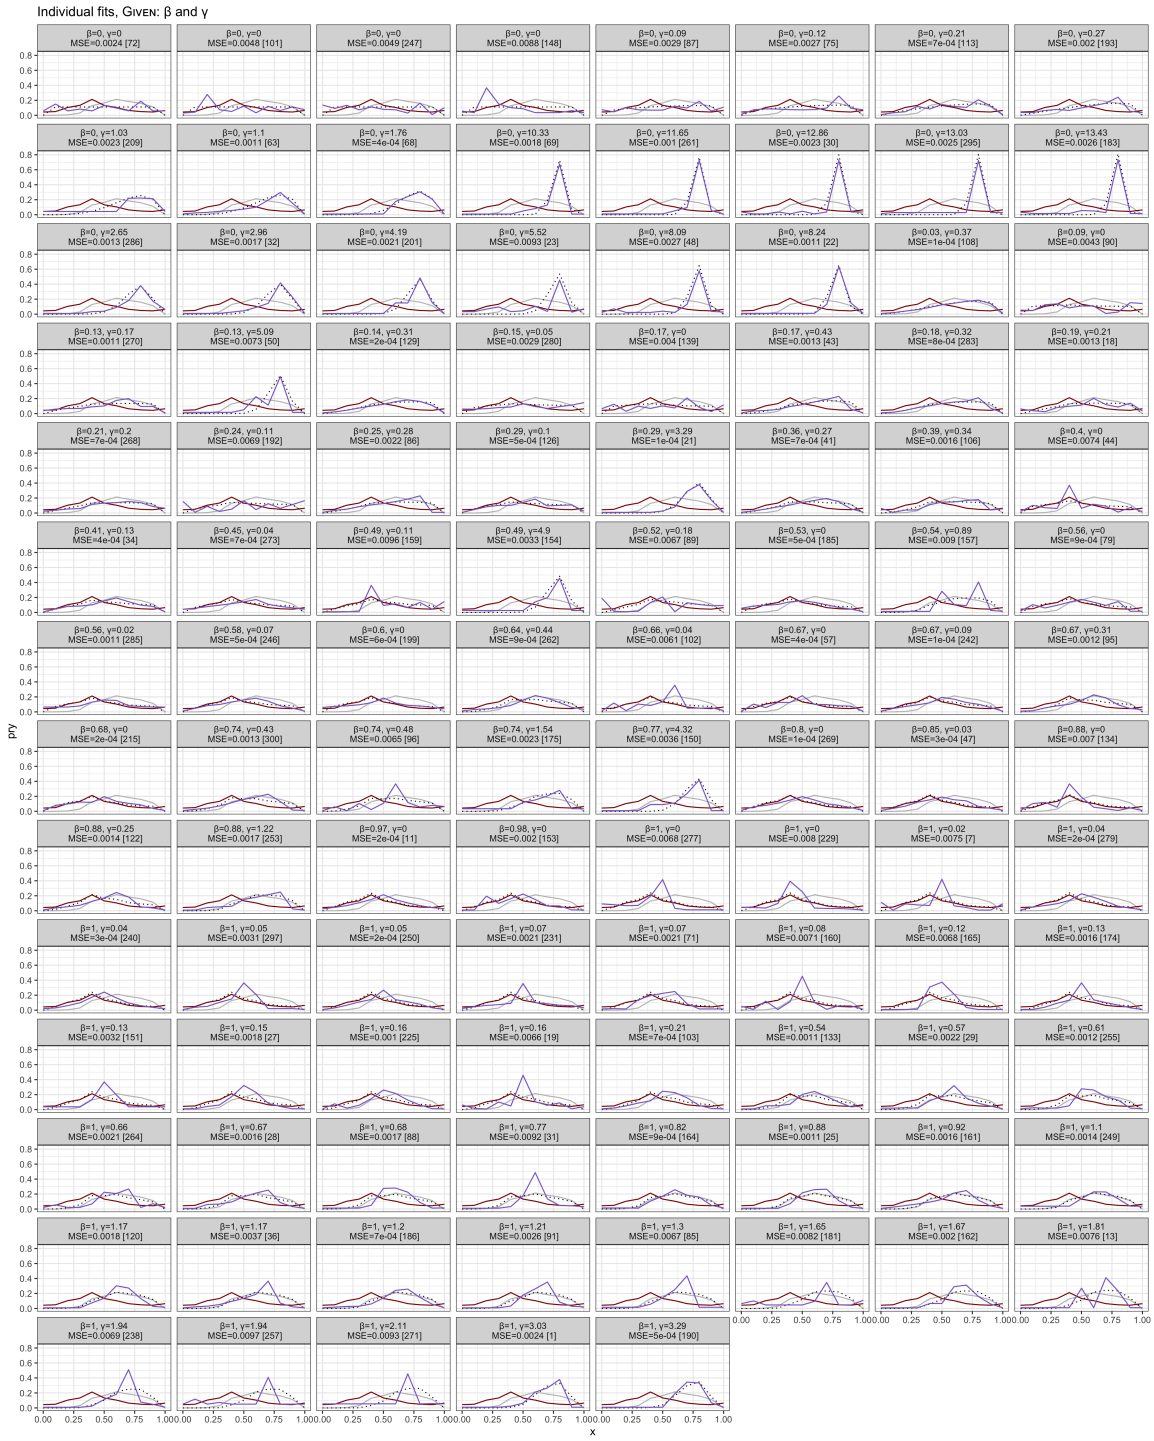

Note. Each good individual participant in the GIVEN condition. The red line is the reported prior, the light purple line is the reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 29

Experiment 3 BAD *participants: Each person in the* ESTIMATED *condition, fitting* β *and* γ

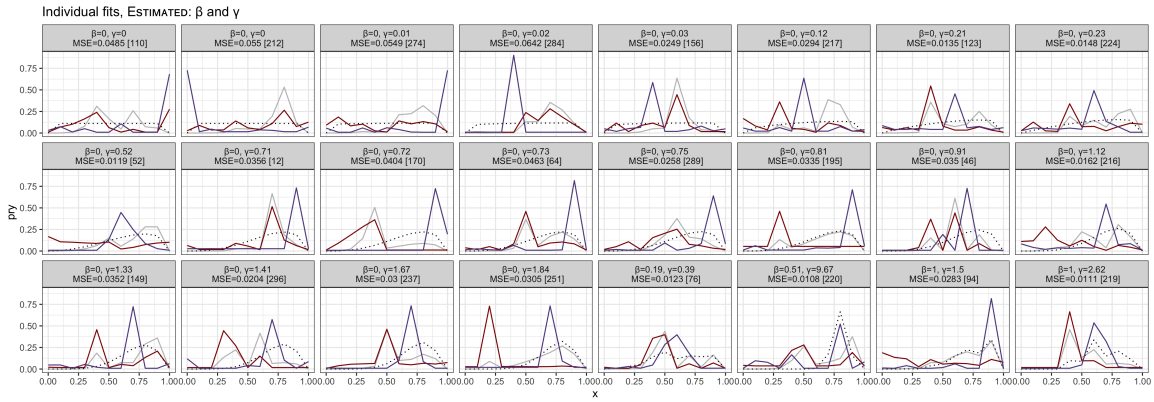

Note. Each bad individual participant in the ESTIMATED condition. The red line is the reported prior, the dark purple line is the reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.

Figure 30

Experiment 3 BAD *participants: Each person in the* GIVEN *condition, fitting* β *and* γ

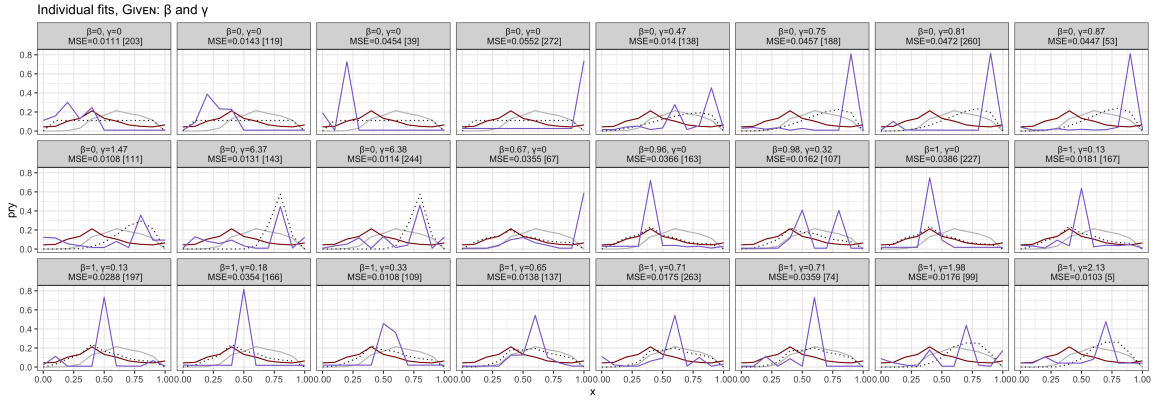

Note. Each bad individual participant in the GIVEN condition. The red line is the reported prior, the light purple line is the reported posterior, the grey line is the posterior obtained by a well-calibrated Bayesian reasoner with that prior and $\beta = \gamma = 1$, and the dotted black line is the posterior obtained by the best-fit values of β and γ for that person. The grey label for each panel contains those values as well as the mean squared error of the fit (MSE, with 0 being perfect). The number in square brackets is the participant ID.