


# meaning of colors 

Master multiverse visualization reference note (not on diagram)

Regular multiverse visualization reference note

Reference note marked as not containing a relevant topic

Reference note marked as containing only saturated topics

General topic category - organized into at least a partial hierarchy

Rough, initial, developing note topic category

## Analysis Category

## Inspection Category

Interpretation/Conclusion Category

Rule 1:
Grey notes are stone -
leave these master notes
in place as a record

## Rule 2:

When you take a note,
make sure to split it fully
so we don't miss
something important.

## Rule 3:

## Left-right distance on the affinity board has meaning. Closeness = similarity

Step 1:
Take a note from a figure frame.

 equal than others.
[orwell, fifl, refi]


## Step 2:

Split the note fully into discrete goals or visualization tasks (as you see it), bolding the focus.

## Step 3:

Place notes on the affinity diagram Almost touching = about the same Separated a little = related Wide space between = different topic

Rule 4:
Notes and clusters can be eorganized by anyone at any time. Move and arrange them whenever it makes sense.

Step 4:
As clusters form, name what they have in common (fewer words are better). Put this name on a differently-colored note above the cluster of notes.

## equality

$$
\begin{aligned}
& \text { more equal than onh onh } \\
& \text { [orwell, fift, refit }
\end{aligned}
$$


"The specification "curve" shows the estimated effect size across all specifications, sorted by magnitude, accompanied below by a "dashboard chart" indicating the operationalizations behind each result (see e.g., Figure 2). This enables readers to visually identify both the variation in effect size across specifications, and its covariation with operationalization decisions. Specification Curve analysis also includes an inferential component, which combines the results from all specifications into a joint statistical test. It assesses whether, in combination, all specifications reject the notion that the effect of interest does not exist."
[simonsohn2015, fig2, ref2]
"Figure 2 reports the descriptive specification curve for the hurricanes example. The top panel depicts estimated effect size, in additional fatalities, of a hurricane having a feminine rather than masculine name. The figure shows that the majority of specifications lead to estimates of the sign predicted by the original authors (feminine hurricanes produce more deaths), though a very small minority of all estimates are statistically significant ( $\mathrm{p}<.05$ ). The point estimates range from -1 to +12 additional deaths."
[simonsohn2015, fig2, ref4]
"The bottom panel of the figure tells us which analytic decisions produce different estimates. For example, we can see that obtaining a negative point estimate requires a fairly idiosyncratic combination of operationalizations: (i) not taking into account the year of the storm, (ii) operationalizing severity of the storm by the log of damages, (iii) conducting an OLS regression, etc. A researcher motivated to show a negative point estimate would be able to report twenty different specifications that do so, but the specification curve shows that a negative point estimate is atypical." [simonsohn2015, fig2, ref5]
"Returning to Figure 1, this appears to be a Panel C situation. Original authors and critics disagree on the set of valid specifications to run. The specification curve results from Figure 2 show that, while such disagreements may be legitimate and profound, we do not need to address them to determine what to make of the hurricanes data. In particular, the figure shows that even keeping the same set of observations as the original study and treating damages in the same way as treated in the original, modifying virtually any arbitrary analytical decision renders the original effect nonsignificant. Readers need not take a position on whether it does or does not make sense to include a damages x pressure interaction in the model to determine if the original findings are robust." [simonsohn2015, fig2, ref6]
"Figure 2 shows that PNAS could have published nearly 1,700 letters showing individual specifications that make the effect go away (without deviating from the original red circle). It also could have published 37 responses with individual specifications showing the robustness of the findings. It would be better to publish a single specification curve in the original paper." [simonsohn2015, fig2, ref7]



"Figure 1. Modeling distribution of union wage premium. Note: Kernel density graph of estimates from 1,024 models. Vertical line indicates the preferred estimate of an 11 percent union wage premium as reported in Table 2." [young2017, fig1, ref1]
"Application 1: The Union Wage Premium. Before proceeding to more detailed aspects of model robustness, we illustrate the basic approachrobustness to the choice of controls-using a data set included in Stata, the 1988 wave of the National Longitudinal Survey of Women. We estimate the effect of union membership on wages (i.e., the union wage premium) controlling for 10 other variables that may be correlated with hourly wages (and union membership; (see Table 2). The coefficient on union, 11.1, means that union members earn about 11 percent more than nonunion members. This is on the low side of conventional estimates, which center around a 15 percent premium (Hirsch 2004)." [young2017, fig1, ref2]
"Next, we report the robustness of this finding to the choice of control variables in the model. Does this finding hinge on sets of control variables, or do the findings hold regardless of what assumptions are made over the control variables? Table 3 shows that there are 1,024 unique combinations of the control variables. Running each of these models and storing all of the estimates, we graph the modeling distribution in Figure 1. The result appears strongly robust. The estimated coefficient on union membership is positive and significant in every possible combination of the control variables: both the sign stability and the significance rate are 100 percent. With this list of possible controls, and using OLS, it is not possible to find an opposite signed or even nonsignificant estimate. Figure 1 shows the modeling distribution as a density graph of all the estimates calculated; the vertical line marks the 11 percent wage premium estimate from Table 2. Estimates as low as 9 percent and as high as over 20 percent are possible in the model space." [young2017, fig1, ref3]
"As shown in Table 3, the average estimate across all of these models is 14.0. This simply represents the average coefficient across all models and is not necessarily the most theoretically defensible. The average sampling standard error is 2.4, and the modeling standard error is 2.5-uncertainty about the estimate derives equally from the data and from the model. The combined total (sampling and modeling) standard error is 3.5.6 The robustness ratio-the mean estimate divided by the total standard error-is 4.05. By the standard of a t-test, this would be considered a strongly robust result, which agrees with the 100 percent sign stability and significance rates. Our conclusion is that, within the scope of these model ingredients, the positive union wage premium is a clear and strongly robust result. This suggests that the decline of unionization in America may well have contributed to middle-class wage stagnation-and not just for male workers (Rosenfeld 2014)." [young2017, fig1, ref4]



## "Figure 3. Modeling distributions for the gender effect under different assumptions." [young2017, fig3, ref1]

Influence Analysis of the Gender Effect in Mortgage Lending: For the mortgage lending analysis, Table 6 shows the influence of control variables on the coefficient of interest (female). The Delta-Beta effect of controls is reported in order of absolute magnitude influence. To aid interpretation, we also report Delta-Beta as a percent change in the estimate from the mean of the modeling distribution (2.29 as in Table 7). Two control variables clearly stand out as most influential: marital status and race. The influence estimate for marriage shows that, all else equal, when controlling for marital status the coefficient on female increases by 2.47 , more than doubling the mean estimate across all models. Controlling for race (with the dummy variable "black") also increases the effect size of gender by 1.91 , a full 83 percent higher than the mean estimate. The other controls have much less impact on the estimate and have little model influence."
[young2017, fig3, ref2]
"In essence, there are two distinct modeling distributions to consider which are plotted in Figure 3. In one set of models, the controls for race and marital status are always excluded but all other controls are allowed in the model space (which gives 128 models). Under these assumptions, the estimates of the gender effect are tightly centered around zero, with an almost even split between positive (52 percent) and negative (48 percent) estimates, none of which are statistically significant. Here, there is no evidence at all for a gender effect. In contrast, the second distribution is defined by the opposite assumption: race and marital status must be in the model, but all combinations of the other controls are possible. Under these assumptions, the estimates cluster around a 4.5 percent higher mortgage acceptance rate for women. Both the significance rate and the sign stability are 100 percent- complete robustness. In order to draw robust conclusions from these data, one must make a substantive judgment about two key modeling assumptions: the inclusion of race and marital status. None of the other model ingredients affect the basic conclusion. These two model assumptions determine the results."
[young2017, fig3, ref3]

"In Table 7 , we show our main analysis. Model 1 includes just the base
populations of the origin and destination states and the income tax differences
between them. When the income tax rate in the origin state is higher, there
tends to be more migration from the origin state to other (lower tax)
destinations. Migration flows are 1.4 percent higher for each percentage point
difference in income tax, but the estimate is not statistically significant. Model 2
adds in controls for contiguity, distance, the sales and property tax rates, state
income, and a measure of natural amenities (topographical/landscape
variability). The etax effect is now larger and statistically significant. For each one
point difference in the tax rate, migration flows are 2.4 percent higher. Finally, in
model 3, when using an IRS migration data with the same set of controls, we find
a similar significant effect. This gives seemingly compelling evidence that high
income taxes cause migration to lower tax states."
[young2017, fig4, ref2]
"What this fails to show, however, is the extreme model dependence in this conclusion. Models 2 and 3 are knifeedge specifications, carefully selected to report statistically significant results, and remarkably unrepresentative of the overall modeling distribution. Both models are highly sensitive to adding or deleting insignificant controls, and this set of controls is the only combination among many thousands that yields a significant result in both the ACS and IRS data.
[young2017, fig4, ref3]
"We embrace a wide robustness analysis that relaxes assumptions about possible controls, possible data sources for migration, and alternative estimation commands. There are two controls that we see as absolutely critical to the gravity model: base populations of the origin and destination states. Combinatorially including or excluding these variables produces models that we regard as nonsense, so we impose the assumption that they must be in all models. However, we leave as debatable the controls for distance, contiguity,
other tax rates, economic performance of the states, and a rich set of natural other tax rates, economic performance of the states, and a rich set of natural
amenities which have been previously shown to influenc marat amenities which have been previously shown to influence migration (McGranahan 1999). All possible combinations of these controls give 4,096 models. Moreover, we test (hese models across the two alternative data sets for strategies (Poisson, negative binomial, and OLS log-linear). For each data set, there are three possible estimation commands, and for each (data set $X$
 robustness [young2017, fig4, ref4]
"As shown in Table 8, the tax coefficient is statistically significant in only 1.5 percent of all models. The mean estimate is almost exactly zero, and estimates are evenly split between positive tax flight estimates ( 48.9 percent) and wrong signed negative estimates (51.1 percent). Among the few statistically significan results, the great majority are wrong signed: estimates with negative signs indicate migration toward higher tax states. Only 0.2 percent of estimates are significantly positive compared to 1.3 percent that are significant and wrong signed. The robustness ratio-the mean estimate divided by the total standard error-is 0.01 . The modeling distribution is relatively normal: There are no critically important modeling decisions that generate bimodality in the estimates. As shown in Figure 4, the significant estimates reported in Table 7 above are extreme outliers in the modeling distribution. [young2017, fig4, ref5]
"In this case, when the robustness analysis is so overwhelmingly nonsupportive, the influence analysis has less to work with. However, there are a few informative points. Compared to Poisson, the negative binomial and OLS log linear models give less positive estimates. Estimates from themodels using IRS rather than ACS data are more positive. This suggests that the most sus vidence will come from using Poisson with the IRS data (reported as model 3 models with ACS data. Yet, even when we narrow our robustness testing to the most supportive estimator (Poisson) and data set (IRS), there is weak support: while the sign stability is 100 percent, the income tax effect is significant in only percent of those models. 13 By control variables, the sales tax rate, average income, and the property tax rate have the most positive influence-generating more positive estimates of tax flight when these controls are included. (Note, however, that none of these controls were significant in model 3.) All other controls push the tax migration estimate toward a zero or wrong-signed result, and virtually must be excluded to support the hypothesis." [young2017, fig4, ref6]

"While it is possible to support the tax flight hypothesis with a few knife-edge model specifications, there is remarkably little support even in a more narrow and supportive robustness analysis. This shows how extreme the difference can e between a curated selection of regression results (Table 7) and a rigorous robustness analysis (Table 8). While one offers an existence proof that a significant result can be found, the weight of the evidence frommany srediblemodels gives scant support to the tax migration hypothesis. It remains technically possible that the one-in-a-thousand specifications of Table 7 present the best, most theoretically compelling estimates. If so, authorswould need to carefully explain to readers why such painstakingly exact model assumptions are equired, and why virtually any departure from model 2 or 3 fails to support the conclusions."
[young2017, fig4, ref8]


"Fig. 2. Volcano plots visualizing the vibration of effects (VOE) for four examples, (A) serum vitamin D, (B) serum thyroxine, (C) urinary creatinine, (D) serum a-tocopherol. Two dimensional histogram representation in upper panel and contour scatter plot is in lower panel. All effects are for 1SD change in logged level of variable interest." [patel2015, fig2, ref1]

The third pattern, as exemplified by an indicator of kidney function, urinary creatinine, and mortality, shows an opposite trend (Fig. 2C). For $\mathrm{k}=5-13$ number of adjustment variables, the association tends to become stronger in HR and statistical significance. in HR and statistical significance; however, the trend is less clear for $\mathrm{k}=$ $0-4$, where the median P -values increase. Twenty-six (6\%) of the 417 variables exhibited similar behavior to urinary creatinine where the effect sizes increased and $P$-value decreased for larger k . [patel2015, fig2, ref4]
4.2. Prototypical patterns of the VoE: We describe four prototypical patterns from the set of 417 variables (Fig. 2, see Appendix at www.jclinepi.com for all 417 variables). The first pattern is exemplified by the association between serum levels of vitamin D and mortality (Fig. 2A). All the HR estimates are <1.00, indicating that higher levels of vitamin D tend to be associated with longer survival (all HR <0.76) however, the magnitude of the estimated effeng is eve (he number of adjus eplenuate when adjusting for more variables, from $=0.64$ with no attenuated when adjustment ( $k=0$ ) to 0.75 with all 13 adjustment variables included ( $k=13$ ). In contrast, the P -values are less than the nominal level of statistical significance ( $P$ $=0.05$, black line). Most of the results are centered on HR $\sim 0.72$ and $P \sim 10 \wedge-4$ (two-dimensional mode). In this first pattern, one concludes that although adjustment weakens the magnitude relationship between vitamin D levels and mortality, inferences regarding the relationship are similar throughout all scenarios of adjustment. Of the 417 variables, 53 ( $13 \%$ ) exhibited similar behavior to vitamin D , where all associations were beyond the level of nominal statistical significance, but the association was attenuated with a greater number of adjustment variables (see Fig. S1/Appendix at www.jclinepi.com)."
[patel2015, fig2, ref2]
"The second pattern is exemplified by the relationship between thyroxine and mortality, displays how increasing adjustment might change inference (Fig. 2B). Higher thyroxine levels tend to be associated with longer survival but P-values become greater than the nominal level of statistical
significance ( $P=0.05$ ) with nine or more adjustment variables on average. Of the 417 variables, 91 (22\%) variables had similar behavior to thyroxine in which HR were attenuated and the $P$-values rose above the nominal level of significance ( $P>0.05$ ) as the number of adjusting variables, $k$, increased (see Fig. S1 and Table S3/ Appendix at www.jclinepi.com)." [patel2015, fig2, ref3]
"In the last pattern, as exemplified by a-tocopherol (Fig. 2D), the estimated HRs can be both greater and less than the null value (HR > 1 and $H R<1$ ) depending on what adjustments were made. We call this the Janus effect after the two-headed representation of the ancient Roman god. For atocopherol, most of the HR and P -values were concentrated around 1 and nonsignificance, respectively. However, $1 \%$ of the models had an HR < 0.875 ( $12.5 \%$ decreased risk of death for 1SD increase in exposure) with a nominally significant $P$ value ( $\mathrm{P}<0.05$ ), whereas another $1 \%$ of the models had HR > 1.05 (5\% increased risk for death for 1SD increase of exposure), albeit without reaching nominal significance. The Janus effect is common: 131 (31\%) of the 417 variables had their 99th percentile $H R>1$ and their 1st percentile $H R<1$." [patel2015, fig2, ref5]
'Examples such as those in Fig. 2A-Drepresented heVot patterns for $72 \%$ of the 417 associations. Other patterns included VoE where all P-values were 0.05 and the strength of the association decreased ( $n=50,12 \%$ ), increased ( $n=27,6 \%$ ), or showed no dependence ( $n=15,4 \%$ ) with increasing number of adjustment variables $k$ (see Table S3/Appendix at www.jclinepi.com). Rarer patterns included variables where all P-values were $<0.05$ and there was an increasing strength of association ( $n=5,1 \%$ ) or no lear relationship with increasing $k(n=4,1 \%)$, and those having $P$-values with a range less than and greater than 0.05 with no clear relationship with k ( n = 15, 4\%)."
[patel2015, fig2, ref6]

"Fig. 3. Volcano plots visualizing vibration of effects
(VoE) for three examples with multiple "modes." (A) (VoE) for three examples with multiple "modes." (A)
The 2D histogram for 1SD increase of the logarithm of serum cadmium, (B) volcano scatter plot with of serum cadmium if smoking was included in the model (yellow) or smoking not included in model (black). (C) Volcano scatter plot for serum cadmium models with drink five per day (yellow) or models without drink five per day (black). (D) The 2D histogram for ISD increase of the logarithm of serum triglycerides, (E) with total cholesterol
 ot included in model (black). (F) With any diabetes yellow) or models wit [patel2015, fig3, ref1]
"To identify the key variable(s) that separated these different distributions, we visualized the VoE by coloring each point on whether it included (or did not seprate visualizations. In serum cadmium we observed the two distinct mod separate visualizations in me observed the two distinct modes smoking (Fig 3B) Specifically, models that contained the smoking adjustment smoking (Fig. 3B). Splin points) had HR lower the smoking adjustment and lower-log10(P-values) (Fig 3B, black points) One source of cadmium exposure includes smoking, and we concluded that the source of cadmum exposure and exposure to cadmium might be triving the correlation between smoking and exposure to cadmium might be driving th
multimodal behavior of VoE. Furthermore, we observed that models that included (or did not include) alcohol drinking also resulted in separate modes in P-values (Fig. 3C)." [patel2015, fig3, ref3]

We observed three modes in the association between triglyceride levels and mortality (Fig. 3D-F). The multimodal plots indicated that total cholesterol and diabetes were driving these modes. For example, in models that did not contain these two djustments, the associations had smaller P-value and a smaller range of $H R$. Furthermore, in model containing diabetes, HR were attenuated. The multimodal pattern seems reasonable in light of the high correlation between triglyceride levels and total high correlation between triglyceride levels and tota molesterol levels/risk for diabetes. We observed a including fasting blood glucose and insulin (see Fig. including fasting blood glucose and [patel2015, fig3, ref4]
4.3. Identifying "multimodality of effects" with VoE: By empirically estimating the VoE, it is also possible to detect whether one or more adjustment variables make a marked difference in the results, leading to multiple modes (Fig. 3) which we call multimodality of effects. Multimodality of effects was clearly seen in 71 of the $417(17 \%)$ assessed variables. For example, the overall VoE for serum cadmium on example, the overall VoE for serum cadmium on
mortality indicates strong association with mortality (Fig. 3A); all of the HRs are $>1.2$ per 1 SD change in serum cadmium levels, and $P$-values in all analytical scenarios are $<0.05$. However, two modes are visually evident (Fig. 3A)."
[patel2015, fig3, ref2]


D

"Fig. 4. Cumulative distributions of vibration of effects (VoE) for 417 variables. (A) Absolute deviation of HR from 1, (B) $\log 10(P$-value), (C) relative hazard ratio (RHR), (D) relative P-value (RP). Examples shown in Figures 1-3 are shown in the distribution." [patel2015, fig4, ref1]
"4.4. Summary of common patterns of the VoE: Figure 4 shows the distribution of the fold deviation of HR from the null (HR = 1.00), the - $\log 10(P$-value), RHR, and RP for all 417 variables considered. The "fold deviation" is the difference of the median VoE-estimated HR from 1 (the null value). The median fold deviation was 1.13 -fold (25th percentile: 1.05 -fold, 75 th percentile: 1.24 -fold, Fig. 4A). Moreover, $50 \%$ of the variables had a median P -value less than or greater than 0.27 (25th percentile: $0.04,75$ th percentile: 0.59 , Fig. $4 B$ ). The median RHR was 1.15 (5th percentile: 1.07 , 25th percentile: 1.11 , 75 th percentile: 1.22 , 95 th percentile: 1.70, Fig. 4C). The median RP was 1.07 (5th percentile: $0.31,25$ th percentile: $0.589,75$ th percentile: $2.03,95$ th percentile: 5.09 ). We observed that most associations could vary by at least 1.15 -fold in the magnitude of the HR and by one order of magnitude [log10(P-value)] in the level of statistical significance, and much larger changes were not uncommon. We observed a weak correlation between RHR and RP (see Fig. S2/ Appendix at www.jclinepi.com, $\mathrm{p}=0.09, \mathrm{P}=0.06$ )."
[patel2015, fig4, ref2]
"Returning to the prototypical examples that we discussed previously, the RHR for vitamin D and thyroxine was moderate 1.14 (44th percentile) and 1.15 ( 51 st percentile; Figs. $4 C, 2 A$, and $B$ ). However, their RPs were among the largest and equal to 4.7 (93rd percentile) and 2.90 ( 84 th percentile), respectively (Figs. 4D, 2A, and B). For urinary creatinine, the scenarios of adjustment had less prominent VoE. The RHR and RP for urinary creatinine was 1.07 (5th percentile) and 0.98 (47th percentile; Fig. 4C and D)."
[patel2015, fig4, ref3]
"The RHR for a-tocopherol (with the Janus effect) was higher (1.21, 71st percentile, Fig. 4C). Variables that demonstrated multimodality, such as serum cadmium and triglycerides, tended to have larger VoE. For example, serum cadmium had an RHR of 1.29 ( 82 nd percentile) and one of the highest RPs, 8.29 (99th percentile). Serum triglycerides had an RHR of 1.18 (64th percentile) and an RP of 1.93 (73rd percentile)."
[patel2015, fig4, ref4]
arslan2018 - fig4
M_1. Main model (all), BC+BCi-
M_p1. Contrinuous, BC -
"Figure 4. Robustness checks for predictors. Coefficient plot showing a consistent effect of the fertility predictor among naturally cycling women (red) but not hormonal contraception users (black) across several predictor and model specifications (explained in further detail in the text). FC = forward counted from last menstrual onset, BC = backward-counted from observed next menstrual onset, $\mathrm{BCi}=$ backwardcounted from inferred next menstrual onset."
[arslan2018, fig4, ref1]
"In models M_p1 to M_p11, we tested different estimates of the fertile window as our predictor to address the concerns about varying standards described in Methodological issues. We compared all combinations of a narrow window broad window, continuous estimates, and backward- and forwardcounting When we used a continuous fertile window predictor, we also adjusted for premenstrual and menstrual days. We found that including adjustments for menstruation and pre-menstruation (M_c3) reduced effect sizes for the fertile window predictor. We could not always adjust for (pre- )menstruation when using a narrow window predictor because of model convergence problems. After taking this into account, we found no systematic pattern in which certain predictors (narrow or broad window, forward or backward counted) had larger effect sizes than others across outcomes (see Figure 4). However, continuous curves over backward-counted days (Figure 3) matched the predicted pattern more closely than curves over forward-counted days (see supportive website osf.io/pbef2)."
[arslan2018, fig4, ref2]
"Although it is difficult to compute an equivalent of Cohen's d for multilevel models, our comparable effect size estimates ranged from 0.12 to 0.43 . These effect sizes are disattenuated for measurement error in the predictor, but not in the outcome. Some were hence only a quarter of the smallest effect size (0.4) considered in Gangestad et al.'s (2016) simulations and sample size recommendations. Empirically, had we used sample sizes like the studies we were replicating, none of the effects reported here would have been significant. Whether the fertility predictor was formed based on forward- or backwardcounting, narrow, broad, or continuous fertile phases seemed to make less of a difference (Figure 4), except that predictors using more data are preferable and that (pre- )menstruation should be adjusted for. While the absolute sizes of the effects we found were not huge, their practical implications might still be noteworthy. The effects on in-pair desire are, for instance, comparable with reported effects of hormonal contraceptive use on sexual desire in a randomised controlled trial (Zethraeus et al., 2016). Moreover, we found evidence for substantial inter-individual variation, so that effects that are small on average might be substantial for some women."
[arslan2018, fig4, ref3]

| Team | Analytic Approach | Odds Ratio |
| ---: | :--- | :---: |
| 12 | Zero-Inflaped Poisson Regression | 0.89 |
| 17 | Bayesian Logistic Regression | 0.96 |
| 15 | Hierarchical Log-Linear Modeling | 1.02 |
| 10 | Multilevel Regression and Logistic Regression | 1.03 |
| 18 | Hierarchical Bayes Model | 1.10 |
| 31 | Logistic Regression | 1.12 |
| 1 | OLS Regression With Robust Standard Errors, Logistic Regression | 1.18 |
| 4 | Spearman Correlation | 1.21 |
| 14 | WLS Regression With Clustered Standard Errors | 1.21 |
| 11 | Multiple Linear Regression | 1.25 |
| 30 | Clustered Robust Binomial Logistic Regression | 1.28 |
| 6 | Linear Probability Model | 1.28 |
| 26 | Hierarchical Generalized Linear Modeling With Poisson Sampling | 1.30 |
| 3 | Multilevel Logistic Regression Using Bayesian Inference | 1.31 |
| 23 | Mixed-Model Logistic Regression | 1.31 |
| 16 | Hierarchical Poisson Regression | 1.32 |
| 2 | Linear Probability Model, Logistic Regression | 1.34 |
| 5 | Generalized Linear Mixed Models | 1.38 |
| 24 | Multilevel Logistic Regression | 1.38 |
| 28 | Mixed-Effects Logistic Regression | 1.38 |
| 32 | Generalized Linear Models for Binary Data | 1.39 |
| 8 | Negative Binomial Regression With a Log Link | 1.39 |
| 20 | Cross-Classified Multilevel Negative Binomial Model | 1.40 |
| 13 | Poisson Multilevel Modeling | 1.41 |
| 25 | Multilevel Logistic Binomial Regression | 1.42 |
| 9 | Generalized Linear Mixed-Effects Models With a Logit Link | 1.48 |
| 7 | Diricillet-Process Bayesian Clustering | 1.71 |
| 21 | Tobit Regression | 2.88 |
| 27 | Poisson Regression |  |



## Fig. 2. Point estimates (in order of magnitude) and $95 \%$ confidence intervals for the effect of soccer players' skin tone on the number of red cards awarded by referees. Reported results, along with the analytic approach taken, are shown for each of the 29 analytic teams. The teams are ordered so that the smallest reported effect size is at the top and the argest is at the bottom. The asterisks indicate upper bounds that have been truncated to increase the interpretability of the plot; the actual upper bounds of the confidence intervals were 11.47 for Team 21 and 78.66 for Team 27. OLS = ordinary least squares WLS = weighted least squares." <br> [silberzahn2017, fig2, ref1]

"What were the consequences of this variability in analytic approaches? Figure 2 shows each team's estimated effect size, along with its $95 \%$ confidence interval (CI). As this figure and Table 3 show, the estimated effect sizes ranged from 0.89 (slightly negative) to 2.93 (moderately positive) in odds-ratio (OR) units; the median estimate was 1.31. The confidence intervals for many of the estimates overlap, which is expected because they are based on the same data. Twenty teams (69\%) found a significant positive relationship, $\mathrm{p}<.05$, and nine teams (31\%) found a nonsignificant relationship. No team reported a significant negative relationship." [silberzahn2017, fig2, ref2]


"The teams also varied in their approaches to handling the nonindependence of players and referees, and this variability also influenced both median estimates of the effect size and the rates of significant results. In total, 15 teams estimated a fixed effect or variance component for players referees, or both; 12 of these teams reported
significant effects (median $O R=1.32, \mathrm{MAD}=0.12$ ). Eight teams used clustered standard errors, and 4 . these teams found significant effects (median OR = hese teams found significant effects (median OR = $1.28, \mathrm{MAD}=0.13$ ). An additional 5 teams did noter
account for this artifact, and 4 of these teams arcount or this artifact, and 4 of these teams
reported significant effects (median OR $=1.39$, MAD $=0.28$ ). The remaining team used fixed effects for the referee variable and reported a nonsignificant esult ( $O R=0.89$ )."
sillberzahn2017, fig3, ref3]
"What were the results obtained with the differen types of analytic approaches used? Teams that employed logistic or Poisson models tended to report estimates that were larger than those of teams that used linear models (see the effect sizes in Fig. 3, in which the teams are clustered according to the distribution used for analyses). Fifteen teams used logistic models, and 11 of these teams found a significant effect (median $O R=1.34$; median absolution deviation, or MAD $=0.07$ ). Six teams used significant effect (median $\mathrm{OR}=1.36, \mathrm{MAD}=0.08$ significant effect (median OR $=1.36, \mathrm{MAD}=0.08$ ). Of significant effect (median $\mathrm{OR}=1.21, \mathrm{MAD}=0.05$ ) signicanterect (median os lassified as
final 2 teams used models final 2 teams used models classified as significant effects (ORs $=1.71$ and 2.88 , respectively)." [silberzahn2017, fig3, ref2]
silberzahn2017-tab4

Table 4. Covariates used by each team. Team numbers are listed on the top and covariates on the left. A shaded box indicates that the corresponding team used the covariate in their final model. The table is ordered by the frequency by which each covariate was used
[silberzahn2017, tab4, ref1]

Table 4. Covariates used by each team. Team numbers are listed on the top and covariates on the left. A shaded box indicates that the corresponding team used the covariate in their final model. The table is ordered by the frequency by which each covariate was used.

Twenty-nine independent teams of researchers submitte analytical approaches and refined these throughout the crowdsourcing project. Table 2 shows each team's final analytic technique, model specifications and reported effect size. 3 Analytic techniques ranged from simple linear regression to complex multilevel regression and Bayesian approaches. Teams also varied highly in their decisions regarding which covariates to include (see R7.1). Table 4 shows that the 29 teams used 21 unique combinations of covariates. Apart from the variable 'games', which was used by all teams, just one covariate (player position, $62 \%$ ) was used in more than half of the analytic
strategies and three were used in just one analysis. Two sets of covariates were used by three teams each, and four sets of covariates were used by two teams each. The remaining 15 teams used a unique combination of covariates.
[silberzahn2017, tab4, ref2]
physical touch
physical no touch
virtual prop
virtual mouse
"Figure 2: Excerpt from the
mini-paper Freqentist, showing widgets embedded in the text in Bret Victor's [94] style. Operating a widget style. Operating a widget
changes one aspect of the changes one aspect of the
analysis and immediately analysis and immedia
updates the figure." updates the figure."
[dragicevic2019, fig2, ref1]

# 0 s 10 s 20 s 30 s 40 s 50 s 60 s 70 s 

Figure 3. Average task completion time (geometric mean) for each condition. Error bars are $\mathbf{9 5 \%}$ t-based CIs.

We focus our analysis on task completion times, reported in Figures 3 and 4. Dots indicate sample means, while error bars are $95 \%$ confidence intervals computed on logtransformed data [6] using the t-distribution method. Strictly speaking, all we can assert about each interval is that it comes from a nrocedure desioned to canture the
"The Freqentist example [36] is a reanalysis of a CHI
study evaluating physical visualizations [51]. It is malla ideas for a typical frequentist analysis with ideas for a typical frequis). The alysis with onfidence intervals (Cls). The results of the analysis are initially identical to the original paper, including the two figures reporting mean task completion time per technique and pairwise comparisons, with $95 \%$ Cls. Four aspects of the analysis can be changed by the reader, which has the effect of immediately updating the two plots and some text elements such as explanations and figure captions. Changes are made by clicking or dragging the elements of the text in blue as in Bret Victor's explorable explanations [94] (see Figure 2)."
explanations [94] (see Figur
[dragicevic2019, fig2, ref2]
"Clicking the "transformed data" text toggles the text to "untransformed data" and updates the two figures with results from the corresponding analysis Although some researchers recommend that completion times be log-transformed [79], other researchers may be suspicious of, or unfamiliar with data transformations-this option reassures them that the results hold for untransformed data. that the results hold for untransformed data.
Similarly, clicking on "tdistribution" switches the text Similarly, clicking on "tdistribution" switches the text
to "BCa bootstrap" and shows the results of the to "BCa bootstrap" and shows the results of the
analysis using non-parametric bootstrap Cls, which analysis using non-parametric bootstrap Cls,
tend to be liberal (i.e., too narrow) with small samples but do not require distributional
assumptions [59]."
[dragicevic2019, fig2, ref4]
"Finally, the plot with the three planned pairwise comparisons (not shown in Figure 2) shows uncorrected Cls, but the reader can apply a Bonferroni correction by clicking on the text "not corrected for multiplicity". Correction for multiplicity is strongly recommended by many but it is not without drawbacks: there is a controversial and complex literature on the topic [31]. To help the reader interpret the CIs correctly, the mini-paper reader interpret the CIs correctly, the mini-paper
contains a paragraph that gives the individual and contains a paragraph that gives the individual and
the family-wise Cl coverage and false positive rates, the family-wise Cl coverage and false positive rates,
which are updated whenever Bonferroni correction which are updated whenever Bonferroni correce
is turned on or off, or whenever the confidence level is changed. More details can be found in the minipaper itself [36]."
[dragicevic2019, fig2, ref5]
"First, horizontally dragging the "95\%" text has the effect of changing the confidence level ( 7 levels are provided from $50 \%$ to $99.9 \%$ ) and updating the length of error bars in the two figures. This allows the reader to appreciate that the $95 \%$ level is arbitrary [66] and thus that Cls should not be interpreted in a strictly dichotomous manner [29]. Meanwhile, readers who insist on interpreting effects as significant or non-significant have the option of changing the customary cutoff of $\alpha=.05$ ( $95 \% \mathrm{Cls}$ ), for example to the $\mathrm{a}=.005$ ( $99.5 \% \mathrm{Cls}$ ) criterion now advocated by some methodologists [15]."
[dragicevic2019, fig2, ref3]
"The Freqentist mini-paper covers a total of $7 \times 2 \times 2 \times 2=56$ unique analyses The paper concludes that the findings from the original study (i.e., good
evidence of a difference for the first
two comparisons, inconclusive results for the third one) are reasonably robust, as they hold across the submultiverse where the confidence level is at 95\% or less."
[dragicevic2019, fig2, ref6]
dragicevic 2019 - fig3

estimate type
-- Hodges-Lehmann
$\xrightarrow{-}$ mean

95\% interval contains 0
-- FALSE
-- TRUE
"Figure 3: Plot from the mini-paper Likert, summarizing point estimates and $95 \% \mathrm{Cls}$ for an effect measured across 4 different experiments (columns) and analyzed using 9 different methods (rows). Clicking on a row label updates the method section. Here no matter how the data are analyzed, no conclusive effect is found for the first three experiments (blue intervals), while there is convincing evidence for an effect in the fourth (red intervals)." [dragicevic2019, fig3, ref1]
"The Likert mini-paper reanalyzes the four experiments in the original InfoVis study [35] using nine different methods covering a broad range of approaches, including parametric vs. non-parametric and frequentist vs. Bayesian. In contrast with the previous mini-paper, all analysis outcomes are summarized in a static overview figure to facilitate comparison. Seven of the nine methods yield simple effect sizes (e.g., mean differences) which are summarized in the plot shown in Figure 3, while the remaining two methods yield log-odds ratios, reported in a different plot (not shown here). By default, the method section in the mini-paper only details the bootstrap method, which was used in the original study. However, clicking on a row label in the figure changes the method section to provide a description and justification of the selected method, an interpretation of its results, and the $p$-value for the fourth experiment (when available)." [dragicevic2019, fig3, ref2]
"The Likert mini-paper covers a total of 9 unique analyses. It concludes that the results are consistent across analyses: no matter how the Likert data are analyzed, no conclusive effect is found for the first three experiments (blue intervals in Figure 3), while there is convincing evidence for an effect in the fourth (red intervals). The results differ slightly nevertheless, and the reader can observe which types of analysis are more conservative and which ones are more liberal."
[dragicevic2019, fig3, ref3]

## Fertility

The classification of women into a high or low fertility group based on cycle day can be done in several ways:

区 Participants with cycle days ranging from 7 to 14 are assigned to the high fertility group, whereas participants with cycle days ranging from 17 to 25 are assigned to the low fertility group [2],
$\square$ days 6-14 are used for high fertility, whereas days 17-27 are used for low fertility [4],
$\square$ days $9-17$ for high fertility and 18-25 for low fertility [5],
$\square$ days $8-14$ for high fertility and 1-7 and 15-28 for low fertility [6], and
$\square$ days $9-17$ for high fertility and 1-8 and 18-28 for low fertility [7].
"The "Constructing the data multiverse" section in Steegen et al. [87] goes through each data processing choice made in the original study [38] and describes alternative choices that could have been reasonably made. The Dataverse minipaper essentially reproduces this section with the difference that the reader can select particular choices. The mini-paper first lists five ways of dichotomizing a particular dependent variable, and lets the reader choose one of them (Figure 4). Four other data processing operations are described afterwards, each with two to three options to choose from. The mini-paper ends with figure showing the result of the selected analysis in the gure showing the result of the selected analysis in the different option is chosen in the text. [dragicevic2019, fig4, ref2]

The Dataverse mini-paper covers $5 \times 2 \times 3 \times 3 \times 2=180$ unique analyses. Steegen et al. [87] summarizes the multiverse by plotting the 180 corresponding $p$-values. While this summary provides an extremely useful overview clearly showing that the original findings are not robust, it does not allow the reader to examine detailed outcomes of specific analyses of interest By making it possible to select any particular analysis and see the
 mor more complete results than a simple summary of $p$ values. As in the Freqentist mini-paper the multiverse be animated, giving a striking demonstration of the variability of effect sizes across the multiverse that can usefully complement the $p$-value summary. [dragicevic2019, fig4, ref3]

- Skeptical 50\%-50\% Optimistic
- Narrow $50 \%-50 \%$ Wide


Region of Practical Equivalence = [-3.4868, 3.4868]

Prior density $\qquad$
"Figure 5: Excerpt from the minipaper Prior depicting the prior and posterior densities. Readers can use the 2 D selection widget (left inset gray box) or drag the highlighted percentages to change the prior." [dragicevic2019, fig5, ref1]
"Unlike other examples, these two axes are continuous. The reader can change their prior either by clicking and dragging on a point in a 2 -dimensional space (see Figure 5), or by clicking and dragging on text sliders (like how confidence level can be adjusted in the Freqentist mini-paper)."
[dragicevic2019, fig5, ref2]
"In the browser, as users interact with Tangle widgets or our 2D widget (Figure 5) to move along the two dimensions (location and scale), we calculate the weights for the prior distributions and the corresponding weights for the posteriors using the above formula. We then calculate the mixture posterior density and visualize it using D3.js in real time."
[dragicevic2019, fig5, ref3]
dragicevic2019-fig6

|  | $\mathrm{r}=0.1$ | $\mathrm{r}=0.3$ | $\mathrm{r}=0.5$ | $\mathrm{r}=0.7$ | $r=0.9$ | Overall |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | pcp-neg | pcp-neg | scatterplot-pos | scatterplot-neg | scatterplot-neg | scatterplot-pos |
|  | scatterplot-pos | scatterplot-pos | pcp-neg | scatterplot-pos | scatterplot-pos | pcp-neg |
|  | scatterplot-neg | scatterplot-neg | scatterplot-neg | pcp-neg | pcp-neg | scatterplot-neg |
|  | stackedbar-neg | stackedbar-neg | stackedbar-neg | stackedbar-neg | ordered line-pos | stackedbar-neg |
|  | ordered line-pos | ordered line-pos | ordered line-pos | ordered line-pos | donut-neg | ordered line-pos |
| ¢ | donut-neg | donut-neg | donut-neg | donut-neg | ordered line-neg | donut-neg |
| © | stackedarea-neg | stackedarea-neg | stackedarea-neg | ordered line-neg | stackedbar-neg | stackedarea-neg |
|  | ordered line-neg | ordered line-neg | ordered line-neg | stackedarea-neg | stackedline-neg | ordered line-neg |
|  | stackedline-neg | stackedline-neg | stackedline-neg | stackedline-neg | stackedarea-neg | stackedline-neg |
|  | pcp-pos | pcp-pos | pcp-pos | pcp-pos | radar-pos | pcp-pos |
|  | radar-pos | radar-pos | radar-pos | radar-pos | pcp-pos | radar-pos |
|  | line-pos | line-pos | line-pos | line-pos | line-pos | line-pos |


|  | $\mathrm{r}=0.1$ | $r=0.3$ | $\mathrm{r}=0.5$ | $r=0.7$ | $r=0.9$ | Overall |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | pcp-neg | pcp-neg | pcp-neg | scatterplot-pos | scatterplot-neg | pcp-neg |
|  | scatterplot-pos | scatterplot-pos | scatterplot-pos | scatterplot-neg | scatterplot-pos | scatterplot-pos |
|  | scatterplot-neg | scatterplot-neg | scatterplot-neg | pcp-neg | pcp-neg | scatterplot-neg |
|  | stackedbar-neg | stackedbar-neg | stackedbar-neg | donut-neg | donut-neg | stackedbar-neg |
|  | donut-neg | donut-neg | donut-neg | ordered line-pos | ordered line-neg | donut-neg |
| $\begin{aligned} & \stackrel{\rightharpoonup}{\omega} \\ & \stackrel{\rightharpoonup}{\circ} \end{aligned}$ | ordered line-pos | ordered line-pos | ordered line-pos | stackedbar-neg | ordered line-pos | ordered line-pos |
|  | stackedarea-neg | stackedarea-neg | stackedarea-neg | ordered line-neg | stackedbar-neg | stackedarea-neg |
|  | stackedline-neg | ordered line-neg | ordered line-neg | stackedarea-neg | stackedarea-neg | ordered line-neg |
|  | ordered line-neg | stackedline-neg | stackedline-neg | stackedline-neg | stackedline-neg | stackedline-neg |
|  | pcp-pos | pcp-pos | pcp-pos | pcp-pos | radar-pos | pcp-pos |
|  | radar-pos | radar-pos | radar-pos | radar-pos | pcp-pos | radar-pos |
|  | line-pos | line-pos | line-pos | line-pos | line-pos | line-pos |

"Figure 6: Left: plot showing a ranking of visualizations in their ability to convey correlation [48]. Right: an alternative plot that could have reasonably come up in an exact replication, created by bootstrapping the experimental dataset. Some results hold (e.g., the bottom of the ranking) while some do not (e.g., the top and middle of the ranking). The mini-paper Dance allows to animate between 100 of those plots." [dragicevic2019, fig6, ref1]
"The mini-paper reproduces the analysis from the original study, with its four plots. It also lets readers replace the original dataset with any of the 100 bootstrap datasets. When the dataset changes, each of the 4 plots changes slightly. More interestingly, animating the multiverse yields a "dance of plots" similar to Cumming's dance of $p$-values [28] and other statistical dances [32], with the difference that the sampling distribution is estimated from data rather than simulated."
[dragicevic2019, fig6, ref2]
"Animating the multiverse of bootstrap datasets allows the reader to appreciate the reliability of the different quantities, trends and patterns depicted by each plot and to carry out "inference by eye" [30]: a pattern that is stable across the multiverse is a good indication that it is reliable. This is an example of the use of hypothetical outcome plots (HOPs) for conveying uncertainty [50, 53]. Compared to static representations of inferential information such as error bars, this technique has the advantage of being applicable to any plot. It is especially useful for revealing statistical uncertainty that is hidden in some plots, such as the ranking plot reproduced in Figure 6. More examples can be found in the minipaper."
[dragicevic2019, fig6, ref3]
decision block, Im model
Fit a linear model with \{\{fixed\}\} terms
decision block, Imer model
Fit a linear mixed model with \{\{fixed\}\}
and \{\{random\}\}terms
Compile $\quad$ universe_1.R
3 universe_2.R
I universe_3.R
I universe_4.R
Universe_213.R
(1) universe_214.R
Universe_215.R
I universe_216.R
Run \& load outputs into


Fig. 1. Authoring and visualizing multiverse analyses with Boba. Users start by annotating a script with analytic decisions (a), from which Boba synthesizes a multiplex of possible analysis variants (b). To interpret the results from all analyses, users start with a graph of analytic decisions (c), where sensitive decisions are highlighted in darker blues. Clicking a decision node allows users to compare point estimates (d, blue dots) and uncertainty distributions (d, gray area) between different alternatives. Users may further drill down to assess the fit quality of individual models (e) by comparing observed data (pink) with model predictions (teal).
[liu2020, fig1, ref1]

To further investigate model quality, Emma drills down to individual universes by clicking a dot in the outcome view. She sees in the model fit view (Fig. 1e) that a model gives largely mismatched predictions.
[liu2020, fig1, ref2]

Clicking a result in the outcome view populates the model fit view with visual predictive checks, which show how well predictions from a given model replicate the empirical distribution of observed data [14], allowing users to further assess model quality (T5). The model fit visualization juxtaposes violin plots of the observed data and model predictions to facilitate comparison of the two distributions (see Fig. 1e). Within the violin plots, we overlay observed and predicted data points as centered density dot plots to help reveal discrepancies in approximation due to kernel density estimation. When the number of observations is large (S1), we plot a representative subset of data, sampled at evenly spaced percentiles, as centered quantile dotplots [25]. As clicking individual universes can be tedious, the model fit view suggests additional universes that have similar point estimates to the selected universe. [liu2020, fig1, ref3]
liu2020 - fig5

(a)

Fig. 5. Decision view and outcome view. (a) The decision view shows analytic decisions as a graph with order and dependencies between them, and highlights more sensitive decisions in darker colors. (b) The outcome view visualizes outputs from all analyses, including individual point estimates and aggregated uncertainty.
[liu2020, fig5, ref1]
[all other references are to existing
categories and tasks]
[liu2020, fig5, [ref2]]


Fig. 6. Facet and Brushing. Clicking a node in the decision view (a) divides the outcome view into a trellis plot (b), answering questions like "does the decision lead to large variations in effect size?" Brushing a region in the outcome view (c) reveals dominant alternatives in the option ratio view (d), answering questions like "what causes negative results?"
[liu2020, fig6, ref1]
[all other references are to existing categories and tasks]
[liu2020, fig6, [ref2]]


Fig. 7. PDFs (a) and CDFs (b) views visualize sampling distributions from individual universes. Toggling these views in a trellis plot allows users to compare the variance between conditions.
[liu2020, fig7, ref1]

Besides aggregated uncertainty, Boba allows users to examine uncertainty from individual universes (Fig. 7). In a dropdown menu, users can switch to view the probability density functions (PDFs) or cumulative distribution functions (CDFs) of all universes. A PDF is a function that maps the value of a random variable to its likelihood, whereas a CDF gives the area under the PDF. In both views, we draw a cubic basis spline for the PDF or CDF per universe, and reduce the opacity of the curves to visually "merge" the curves within the same space. There is again a one-to-one mapping between a visual element and a universe to afford interactions. To help connect point estimates and uncertainty, we draw a strip plot of point estimates beneath each PDFs/CDFs chart (Fig. 7, blue dashes), and show the corresponding sampling distribution PDF when users mouse over a universe in the dot plot.
[liu2020, fig7, ref2]
liu2020 - fig8
 (a)
 (b)

## (b) Removing universes that fail to meet a model quality threshold. liu2020, fig8, ref1]

Boba enables an overview of model fit quality across all universes (T5) by coloring the outcome view with a model quality metric (Fig. 8a). We use normalized root mean squared error (NRMSE) to measure model quality and map NRMSE to a single-hue colormap of blue shades where a darker blue indicates a better fit. [liu2020, fig8, ref3]

Now that Emma understands what decisions lead to null effects, she wonders if these results are from trustworthy models. She changes the color-by field to get an overview of model fit quality (Fig. 8a) and sees that the universes around zero have a poorer fit. She then uses a slider to remove universes that fail to meet a quality threshold (Fig. 8b). [liu2020, fig8, ref2]
liu2020 - fig9


Fig. 9. Inference views. (a) Aggregate plot comparing the possible outcomes of the actual multiverse (blue) and the null distribution (red). (b) Detailed plot showing the individual point estimates and the range between the 2.5th and 97.5th percentile in the null distribution (gray line). Point estimates outside the range are colored in orange. (c) Alternative aggregate plot where a red line marks the expected null effect.
[liu2020, fig9, ref1]

To support users in making inference and judging how reliable the hypothesized effect is (T6), Boba provides an inference view at the end of the analysis workflow, after users have engaged in exploration. Once in the inference view, all earlier views and interactions are inaccessible to avoid multiple comparison problems [60] arising from repeated inference. The inference view contains different plots depending on the outputs from the authoring step, so that users can choose between robust yet computationally-expensive methods and simpler ones. [liu2020, fig9, ref3]

In addition, Boba enables users to propagate concerns in model fit quality to the inference view in two possible ways. The first way employs a model averaging technique called stacking [58] to take a weighted combination of the universes according to their model fit quality. The technique learns a simplex of weights, one for each universe model, via optimization that maximizes the log-posteriordensity of the held-out data points in a k-fold cross validation. Boba then takes a weighted combination of the universe distributions to create the aggregate plot. While stacking provides a principled way to approach model quality, it can be computationally expensive. As an alternative, Boba excludes the universes below the model quality cutoff users provide in Sect. 5.4. The decisions of the cutoff and whether to omit the universes are made before a user enters the inference view.
[liu2020, fig9, ref5]

After an in-depth exploration, Emma proceeds to the final step, asking "given the multiverse, how reliable is the effect?" She confirms a warning dialog to arrive at the inference view (Fig. 9). [liu2020, fig9, ref2]

A more robust inference utilizes the null distribution - the expected distribution of outcomes when the null hypothesis of no effect is true. In this case, the inference view shows an aggregate plot followed by a detailed plot (Fig. 9ab). The aggregate plot (Fig. 9a) compares the null distribution (red) to possible outcomes of the actual multiverse (blue) across sampling and decision variations. The detailed plot (Fig. 9b) shows point estimates (colored dots) against 95\% confidence intervals representing null distributions (gray lines) for each universe. Each point estimate is orange if it is outside the range, or blue otherwise. Underneath both plots, we provide descriptions (supplemental Fig. 1) to guide users in interpretation: For the aggregate plot, we prompt users to compare the distance between the averages of the two densities to the spread. For the detailed plot, we count the number of universes with the point estimate outside its corresponding range. If the null distribution is unavailable, Boba shows a simpler aggregate plot (Fig. 9c) where the expected effect size under the null hypothesis is marked with a red line.
[liu2020, fig9, ref4]


Fig. 10. A case study on how model estimates are robust to contr variables in a mortgage lending dataset. (a) Decision view show
that black and married are two consequential decisions. (b) Ovat black and married are two consequential decisions. (b) with three peaks. (c) Trellis plot of black and married indicates the source of the peaks. (d) Model fit plots show that models produce numeric predictions while observed data is categorical. (e) PDFs of individual sampling distributions show significant overlap of the three peaks.
[liu2020, fig10, ref1]

## The patterns revealed by ad-hoc visualizations in previous work are also readily available in the Boba Visualizer, either in the are also readily available in the Boba Visualizer, either in the

 default views or with two clicks guided by prominent visual cues. The default outcome view (Fig. 10b) shows that the point estimates follow a multimodal distribution with three separate peaks. Clicking the two highlighted (most sensitive) nodes in thedecision view (Fig 10 a) produces a trellis plot (Fig. 10 c) where decision view (Fig. 10a) produces a trellis plot (Fig. 10c), where
each subplot contains only one cluster. From the trellis plot, it is each subplot contains only one cluster. From the trellis plot, it is
evident that the leftmost and rightmost peaks in the overall distribution come from two particular combinations of the influential variables. Alternatively, users might arrive at simila insights by brushing individual clusters in the default outcome view.
[liu2020, fig10, ref3]

We first demonstrate that the default views in the Boba Visualizer afford similar insights on uncertainty, robustness, and decision sensitivity. Upon launching the visualizer, we see a decision graph and an overall outcome distribution (fig. 10). The and married. The outcome view (Fio 10b) shows that the point estimates are highly varied with conflicting implications. The ggregated uncertainty in the outcome view (Fig. 10b background gray area) has a wide spread, suggesting that the possible outcomes are even more varied when taking both sampling and decision variability into account. These observations agree with the summary metrics in previous work hough Boba uses a different, non-parametric method to quantify decision sensitivity, as well as a different method to gregate end-to-end uncertainty.
aggregate end-to-end
[liu2020, fig10, ref2]

Finally, the uncertainty and model fit visualizations in Boba surface potential issues that previous work might have overlooked. First, though the point estimates in Fig. 10b fall into three distinct clusters, the aggregated uncertainty distribution appears unimodal despite a wider spread. The PDF lot (Fig. 10e) shows that sampling distribution from one analysis typically spans the range of multiple peaks, thus xplaining why the aggregated uncertainty is unimodal. Thes by point estimates are not robust when we take sampling by point estina are not robust when we take sampling variations into account. Second, we assess model fit quality by vew (fig. 100). As shown in fig. 10d, while the observed data only takes two possible values, the linear regression model roduces a continuous range of predictions. It is clear from his visual check that an alternative model, for example logistic egression, is more appropriate than the original linear esults with skepticism given the model fit issues. These bservations support our arguments in Sect. 3.2 that ncertainty and model fit are potential blind spots in prior literature.
[liu2020, fig10, ref4]


Fig. 11. A case study on whether hurricanes with more feminine names have caused more deaths. (a) The majority of point estimates suggest a small, positive effect, but there are estimates suggest a small, positive effect, but there are
considerable variations. (b) Faceting and brushing reveal decision combinations that produce large estimates. Coloring by model quality shows that large estimates are from questionable models, qualiy shows that large estimates are from questionable models, and preaicl the the (c) cons and distributions are differe in terms of mode and shape, yet with highly overlapping in terms of
estimates.
[liu2020, fig11, ref1]

But do we have evidence that certain outcomes are less
rustworthy? We toggle the color-by drop-down menu so that each niverse is colored by its model quality metric (Fig. 11b). The large stimates aro by its model quality metric (Fig. 1 lb ). The a poor fit. We further verify the model fit quality by picking example fit. We further verify the model fit quality by picking example universes and examining the model fit view (Fig. 11c). The visual predictive checks confirm issues in model fit, for example the models fail to generate predictions smaller than 3 deaths, while the observed data contains plenty such cases.
[liu2020, fig11, ref2]

Now that we have reasons to be skeptical of the large estimates, the remaining universes still support a small, positive effect. How reliable is the effect? We proceed to the inference view to
compare the possible outcomes in the observed multiverse and the expected distribution under the null hypothesis (Fig. 11d). The two distributions are different in terms of mode and shape, yet they are highly overlapping, which suggests the effect is not reliable. The detail plot depicting individual universes
reliable. The detail plot depicting individual universes
(supplemental Fig. 1) further confirms this observation. Out of the entire multiverse, only 3 universes have point estimates outside the 2.5 th and 97.5 th percentile of the corresponding null distribution.
[liu2020, fig11, ref3]

## liu2020 - tasks

Explicit tasks supported by the Boba visualization system

T1: Decision Overview - gain an overview of the decision space to understand the multiverse and contextualize subsequent tasks.
[liu2020, tasks, T1]
T2: Robustness Overview - gauge the overall robustness of findings obtained through all reasonable
specifications.
[liu2020, tasks, T2]

T3: Decision Impacts - identify what combinations of decisions lead to large variations in outcomes, and what combinations of decisions are critical in obtaining specific outcomes.
[liu2020, tasks, T3]

T5: Model Fit - assess the model fit quality of individual universes to distinguish trustworthy models from questionable ones.
[liu2020, tasks, T5]

T4: Uncertainty - assess the end-to-end uncertainty as well as uncertainty associated with individual universes. [liu2020, tasks, T4]

T6: Inference - perform statistical inference to judge how reliable the hypothesized effect is, while accounting for model quality.
[liu2020, tasks, T6]

```
delgiudice2020 - fig4
```

Figure 4. Results of the full multiverse-style analysis of the simulated dataset. (a) Distribution of $p$-values across 1,216 specifications. (b) Vibration of effects (VoE) plot showing the joint distribution of $p$-values and effect sizes for the same specifications. [delgiudice2020, fig4, ref1]

The distribution of $p$-values and vibration of effects in the full multiverse are shown in Figure 4. The median p was .194. Just $27 \%$ of the effects reached the conventional threshold of $a=.05$ Effect sizes ranged from $b=-.16$ to .25 , with a median of $b=.01$. The VoE plot shows a clear "Janus effect" (see Patel et al., 2015), as the regression coefficients at the 1st and 99th percentiles of the effect size distribution have opposite signs ( -.14 and .21 , respectively). These results could be easily interpreted as indications of poor robustness and replicability. The median effect size across specifications was very close to zero and far from conventional significance thresholds, even though the true effect size in the population was $b=.20$ (before accounting for measurement validity). Investigators using the mean of the multiverse as a "robust" estimate would wrongly conclude that the effect of inflammation on depression is about zero.
[delgiudice2020, fig4, ref2]


Figure 5. Specification curve for the simulated dataset (full multiverse of 1,216 specifications). Blue $=$ positive effect sizes significant at $a=.05$. Red = positive effect sizes significant at $a=.05$. [delgiudice2020, fig5, ref1]


Figure 5 displays a specification curve for the full multiverse. The significant effects are split between positive and negative. The pattern for alternative predictors reflects the impact of measurement validity, which is lower for individual biomarkers (especially with simultaneous entry) and higher for composites. But the central tendency of effects is similar across predictors. As for covariates, inspection of Figure 5 indicates that combinations that include fatigue tend to yield negative effects, whereas the direction tends to be positive when fatigue is excluded. Regardless of the general direction of effects, every combination produces a fair amount of nonsignificant findings. Alternative cutoffs for outliers do not seem to have a systematic impact, except that including all cases shifts the distribution toward somewhat more negative effects.
[delgiudice2020, fig5, ref1]

Model 1 (fatigue as collider)
b



Model 2 (fatigue as mediator)


Figure 6. Results of the principled multiverse-style analyses of the simulated dataset. ( $a, c$ ) Distribution of $p$-values across 6 specifications. (b, d) Vibration of effects (VoE) plots showing the oint distribution of $p$-values and effect sizes for the same pecifications.
delgiudice2020, fig6, ref1

In sum, analyses of the principled multiverses revealed two homogeneous clusters of effects, indicating that the exact biomarker composite employed as a predictor and the choice of cutoff for outliers do not substantially change the conclusions of the study. What does make a difference is whether fatigue is treated as a collider and excluded as a covariate (Model 1), or treated as a mediator and controlled for in the analysis (Model 2). Making an informed decision between these models would require additional empirical evidence (e.g., experimental or quasiexperimental studies), theoretical developments, or both. [delgiudice2020, fig6, ref3]


This is an example of something we don't consider a Composite plot. The axes aren't aligned, and don't appear to have a "super-additive" effect of supporting a task the plots individually can't









|  | outcomespes |
| :---: | :---: |
| $\underline{-}$ |  |
| ㄴ․․․․ | $\underline{\square}$ |
| $\pm$ |  |
| $\equiv$ | ․ |
| ㄴ․ | $\underline{\square}$ |
| $\square$ | 자․ |
| $\underline{\square}$ | W |
| $\underline{\square}$ | $\equiv \square$ |
| $\underline{\underline{2}}$ | $\underline{\square}$ |
|  |  |
|  | = |
| $\underline{\square}$ | $\underline{\square}$ |
| $\underline{\underline{L}} \underline{\underline{L}}$ | $\ldots$ |
|  |  |
|  |  |



