EZ Cognitive Psychometrics:
Bayesian hypothesis testing with the EZ-DDM

Adriana F. Chávez De la Peña, Eunice Shin, & Joachim Vandekerckhove
University of California, Irvine

Introduction

The **EZ-DDM** is a simplified, closed-form version of the three-parameter drift diffusion model that enables rapid parameter estimation from choice and RT summary statistics (i.e., the accuracy rate ($A^{\text{pred}}_{\text{mean}}$), the mean RT ($RT^{\text{pred}}_{\text{mean}}$), and the RT variance ($RT^{\text{pred}}_{\text{var}}$)), through a method of moments [2].

$$\begin{aligned} k &= \exp(-\alpha\nu) \\ A^{\text{pred}}_{\text{mean}} &= \frac{1}{1+k} && (1) \\ RT^{\text{pred}}_{\text{mean}} &= \tau + \left(\frac{\alpha}{2\nu}\right)\left(\frac{k-1}{k+1}\right) && (2) \\ RT^{\text{pred}}_{\text{var}} &= \left(\frac{\alpha}{2\nu^3}\right)\left\{\frac{1-2k\alpha\nu-k^2}{(k+1)^2}\right\} && (3) \end{aligned}$$

In previous work [1], we have shown that casting $A^{\text{pred}}_{\text{mean}}$ as the binomial rate for the correct response count ($A^{\text{obs}}_{\text{total}}$), along with the known sampling distributions of $RT^{\text{obs}}_{\text{mean}}$ and $RT^{\text{obs}}_{\text{var}}$, allow us to build a **hyper-efficient proxy likelihood for the EZ-DDM**, using Equations 1–3.

$$\begin{aligned} A^{\text{obs}}_{\text{total}} &\sim \text{Binomial}\left(A^{\text{pred}}_{\text{mean}}, N\right) && (4) \\ RT^{\text{obs}}_{\text{mean}} &\sim \text{Normal}\left(RT^{\text{pred}}_{\text{mean}}, \frac{RT^{\text{pred}}_{\text{var}}}{N}\right) && (5) \\ RT^{\text{obs}}_{\text{var}} &\sim \text{Normal}\left(RT^{\text{pred}}_{\text{var}}, \frac{2\left[RT^{\text{pred}}_{\text{var}}\right]^2}{N-1}\right). && (6) \end{aligned}$$

This likelihood can support a **hierarchical Bayesian EZ-DDM** that can be easily implemented in any probabilistic programming language and takes only a few seconds to run!

We can extend this with **metaregression structures** to explore covariate effects, for example:

$$\nu_p \sim \text{Normal}(\mu_\nu + \beta x_p, \sigma^2_\nu). \qquad (7)$$

Here, β captures the effect of covariate x_p on the drift rate ν_p, enabling for **Bayesian hypothesis testing** against $\beta = 0$.

Hypothesis testing • Simulation study

- We simulated trial data from a **Wiener** process with a **within-subject t test design** on the drift rate parameter (Eq. 7 with $x_p \in \{0, 1\}$).
- We generated 1,000 data sets for every combination of **number of trials per condition** $T \in \{20, 40, 80, 160, 320\}$ by **number of simulated participants** $P \in \{20, 40, 80, 160\}$ by **true fixed effect size** $\beta \in \{0.0, 0.1, 0.2, 0.4\}$.
- We also added contaminant data.
- We conducted the simulation study across four conditions, defined by a 2×2 factorial design:
  - **Data type:** Clean vs. contaminated.
  - **Summary strategy:** EZ statistics (mean and variance of RTs) vs. Robust statistics (median and IQR-derived variance).

Robust statistics

- EZ-DDM relies on mean and variance of RTs
- Both statistics are sensitive to outliers
- We replace these with 'robust' alternatives:
  - Median (P_{50}) instead of mean
  - Interquartile range approximates variance

$$\sigma^2 \approx \left(\frac{P_{75} - P_{25}}{1.349}\right)^2$$

Data type

- We contaminated 5% of each participant's trials.
- Used a coin flip to choose one of two contamination types: $z_i \sim \text{Bernoulli}(0.5)$:
  - RT noise: Add uniform random noise to the observed RT.
  - Decision noise: Replace the trial with a Wiener draw with $\nu = 0$.

Results

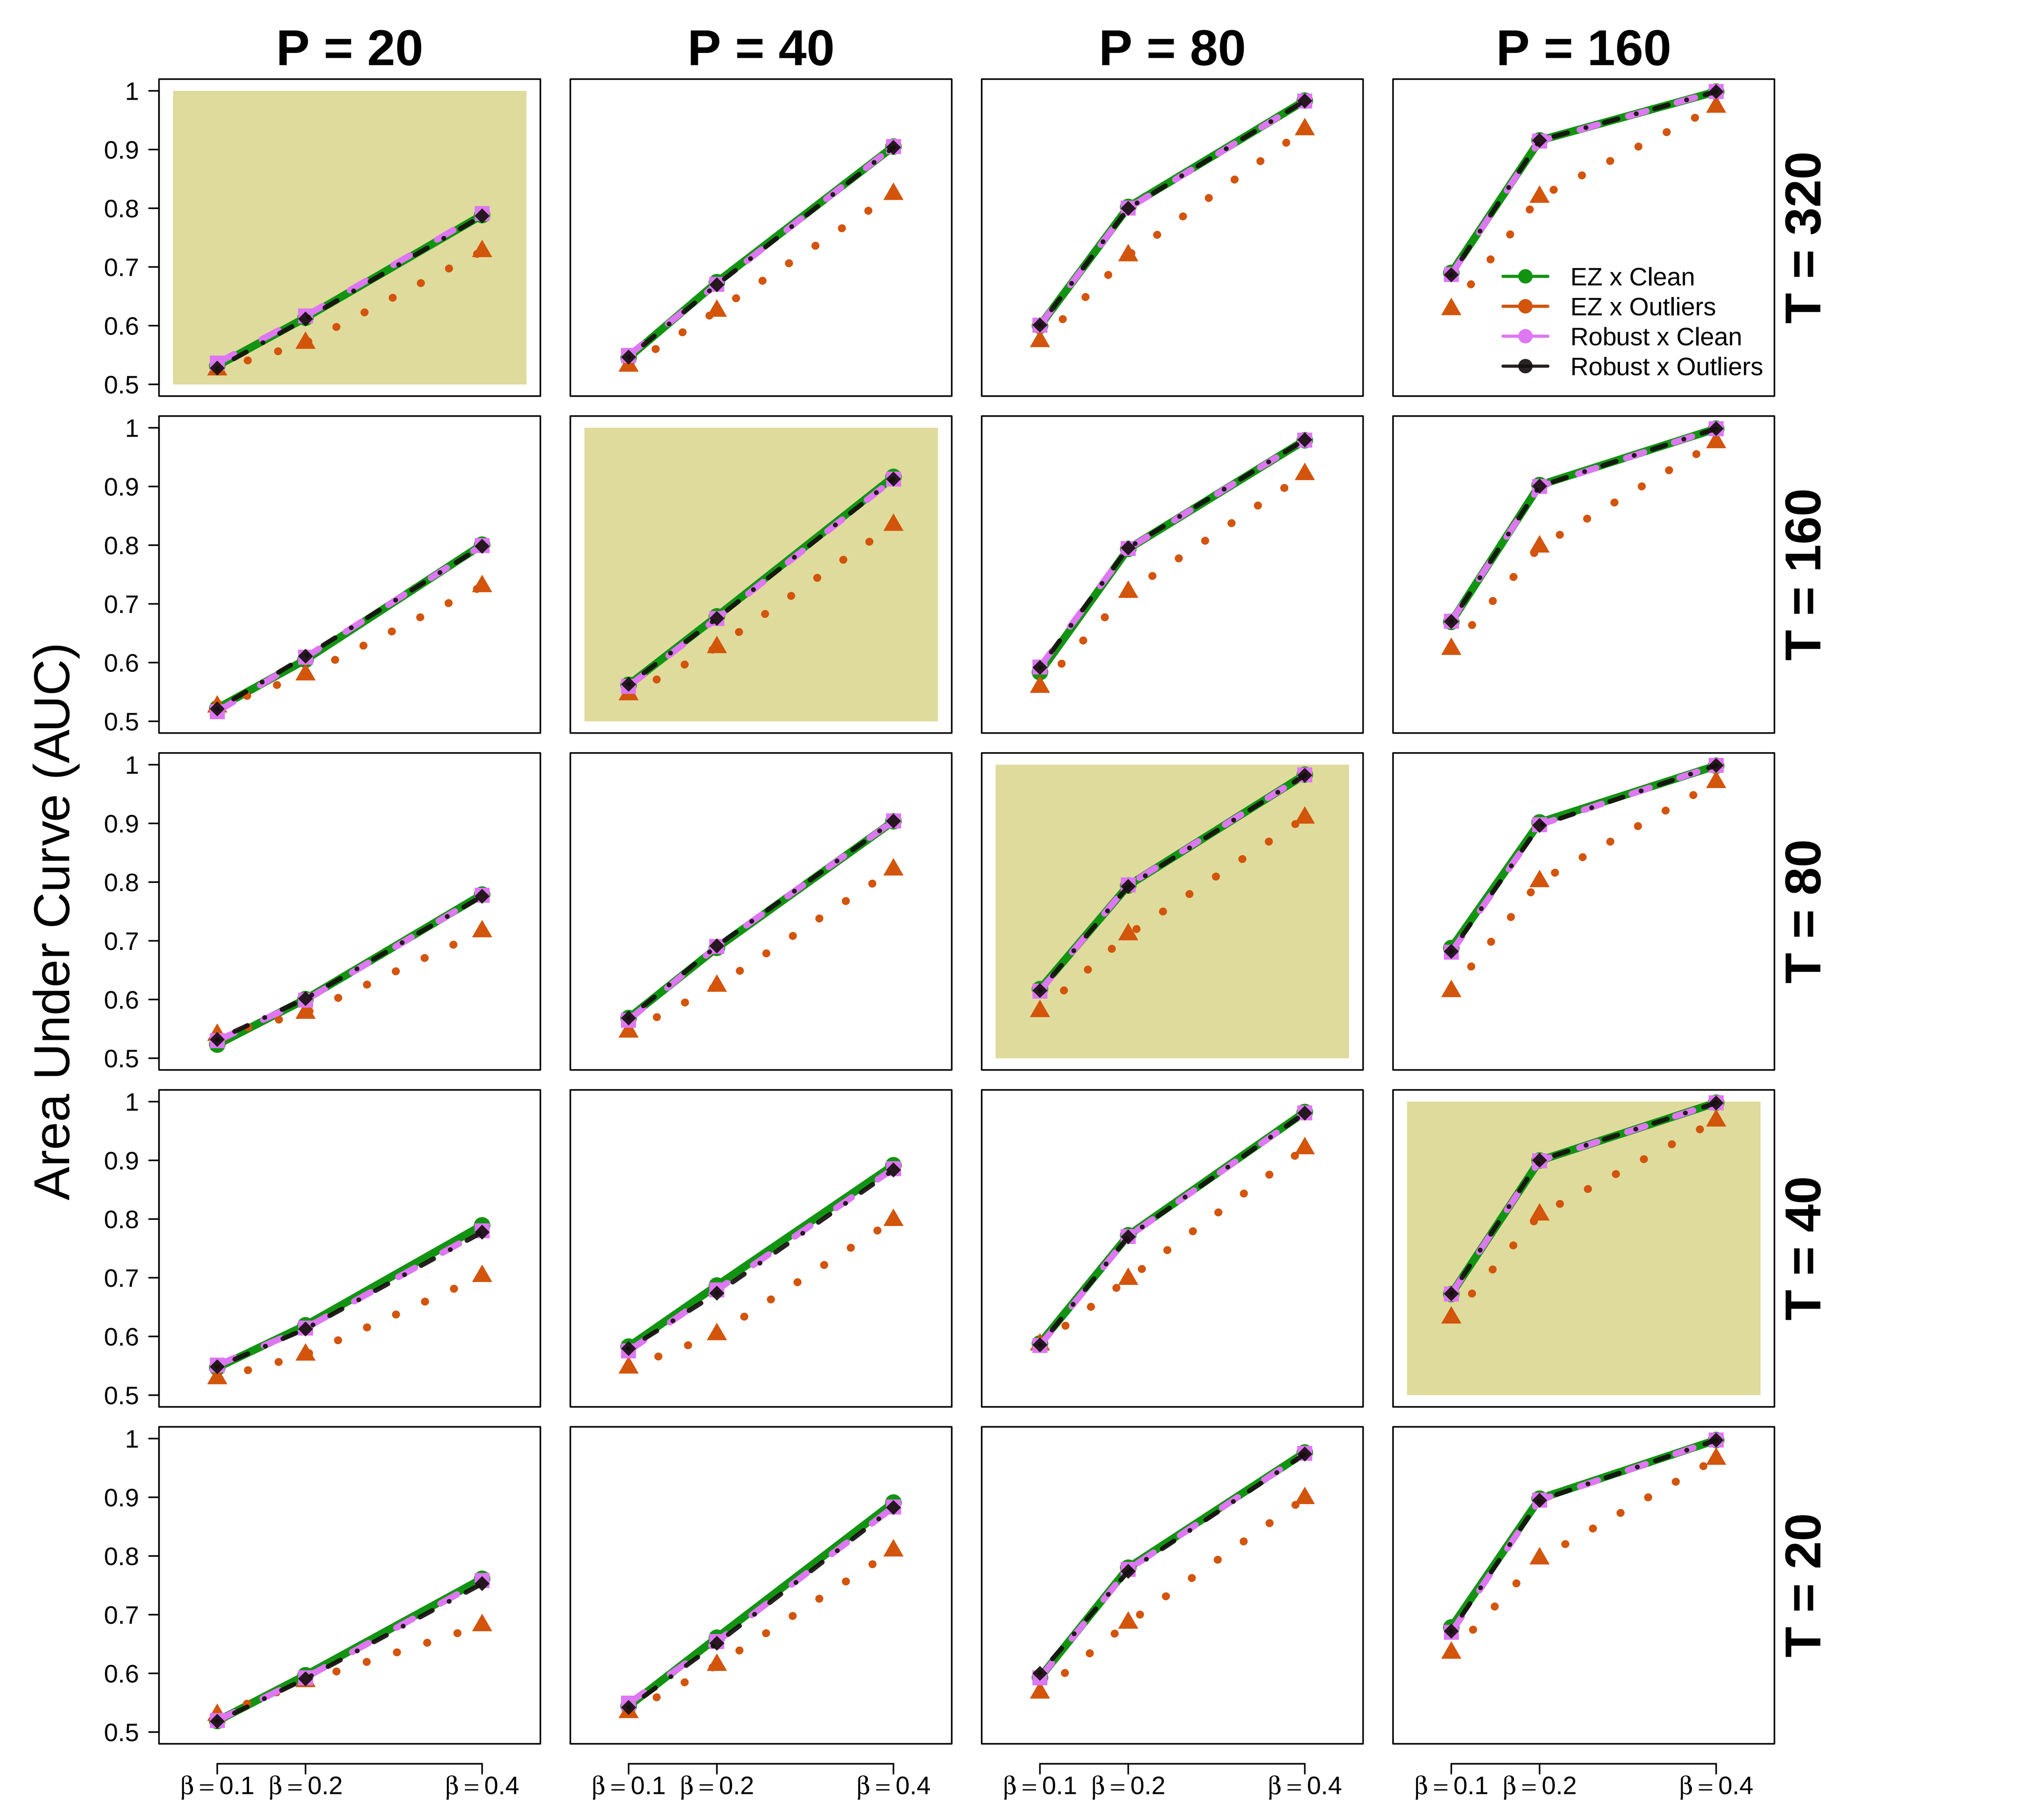

Figure 1: Simulation study results. Each panel shows the AUC (vertical axis) for different true effect sizes (horizontal axis) and simulation conditions (different lines), across different trial and participant sizes (rows and columns).

- Figure 1 shows the Area Under the Curve (AUC) derived from a Receiver Operating Characteristic curve for every fixed effect level β, and simulation condition across all combinations of sample P and trial T sizes.
- The AUC is a metric for the diagnostic accuracy of a Bayes Factor test, quantifying its ability to correctly distinguish between a true effect ($\beta \neq 0$) and a null effect ($\beta = 0$).
- The underlying statistical test is a Bayes Factor calculated with respect to a Region of Practical Equivalence, quantifying the evidence for or against the presence of an effect.

Key findings

- The performance of all methods improves with more observations. However, the number of participants is more important than the number of trials (e.g., all highlighted cells have the same total number of observations).
- The 'Robust' implementation performs similarly to the 'EZ' implementation for clean data.
- The 'EZ' standard implementation is vulnerable to outliers; its AUC drops with contaminated data.
- The 'Robust' implementation is not affected by contamination.

Contact Information

Web cidlab.com
Email achavezd@uci.edu
Github github.com/Adrifelcha/ez-robust

References

[1] Adriana F. Chávez De la Peña and Joachim Vandekerckhove. An EZ Bayesian hierarchical drift diffusion model for response time and accuracy. *Psychonomic Bulletin & Review*, in press.

[2] E.-J. Wagenmakers, H. J. L. van der Maas, and R. P. P. P. Grasman. An EZ–diffusion model for response time and accuracy. *Psychonomic Bulletin & Review*, 14:3–22, 2007.