FAIR DISTRIBUTING

An AI Ethics Case Study

ALGORITHMIC FAIRNESS
WHO GETS ACCESS TO WHAT, AND WHY?

DECISIONS/OPPORTUNITIES OUTSOURCED TO DATA/ALGORITHMS

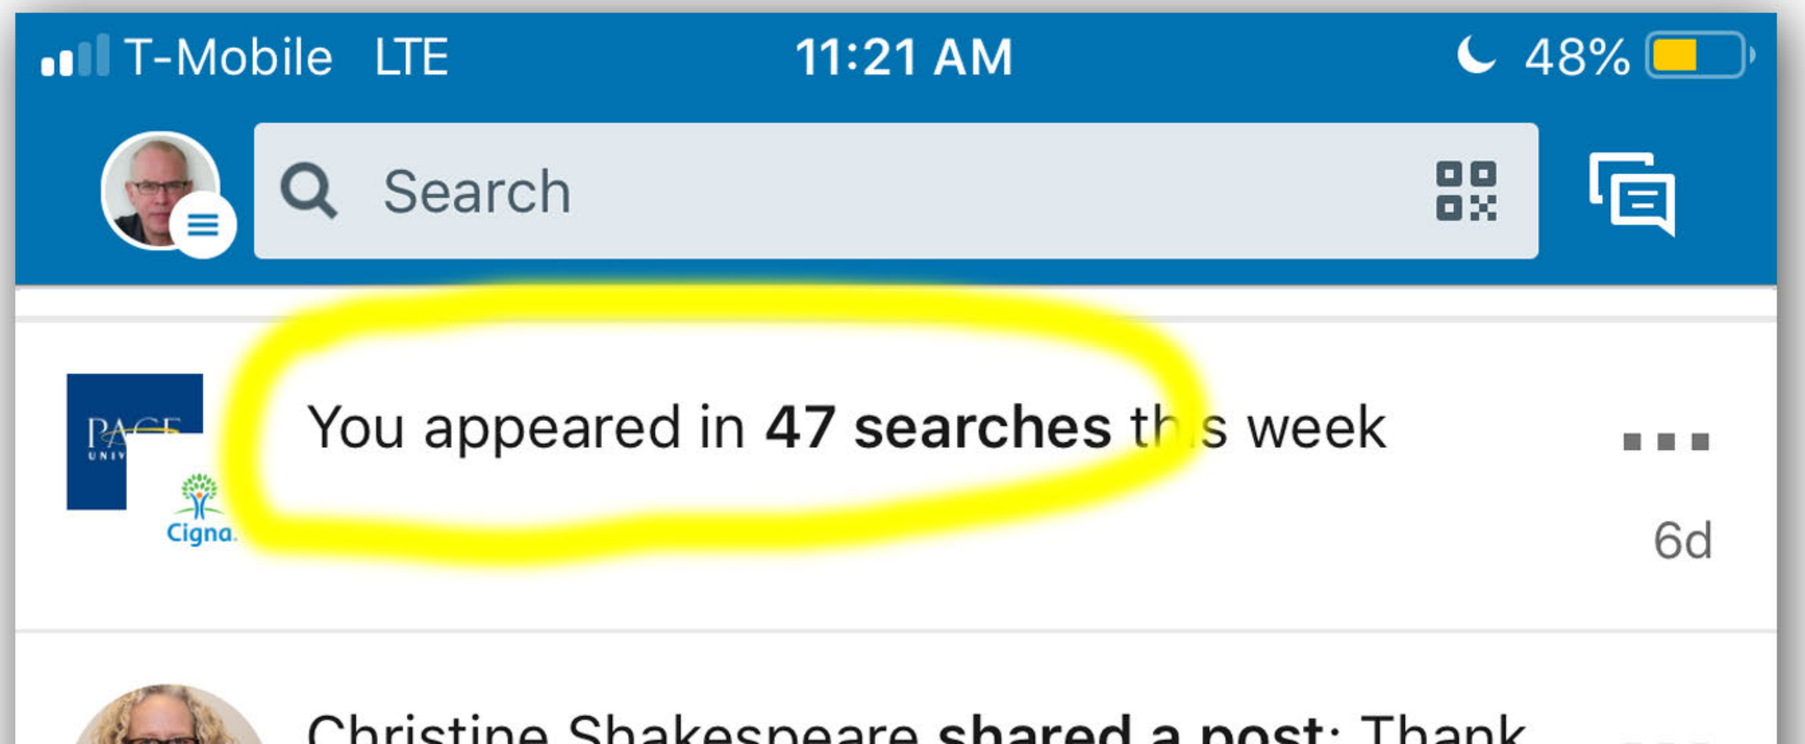

ALGORITHMIC FAIRNESS: FREEDOM DISTRIBUTION

Bernard Parker, Arrested
Rated: High risk of reoffending
Jailed before his trial

Dylan Fugett, Arrested
Rated: Low risk of reoffending
Freed until his trial

WHO GETS ACCESS TO BAIL (FREEDOM)?

ALGORITHMIC FAIRNESS COMPAS

Correctional Offender Management Profiling for Alternative Sanctions

An algorithm estimating pretrial risk of a defendant reoffending. Data includes: prior convictions, age, family situation, etc.

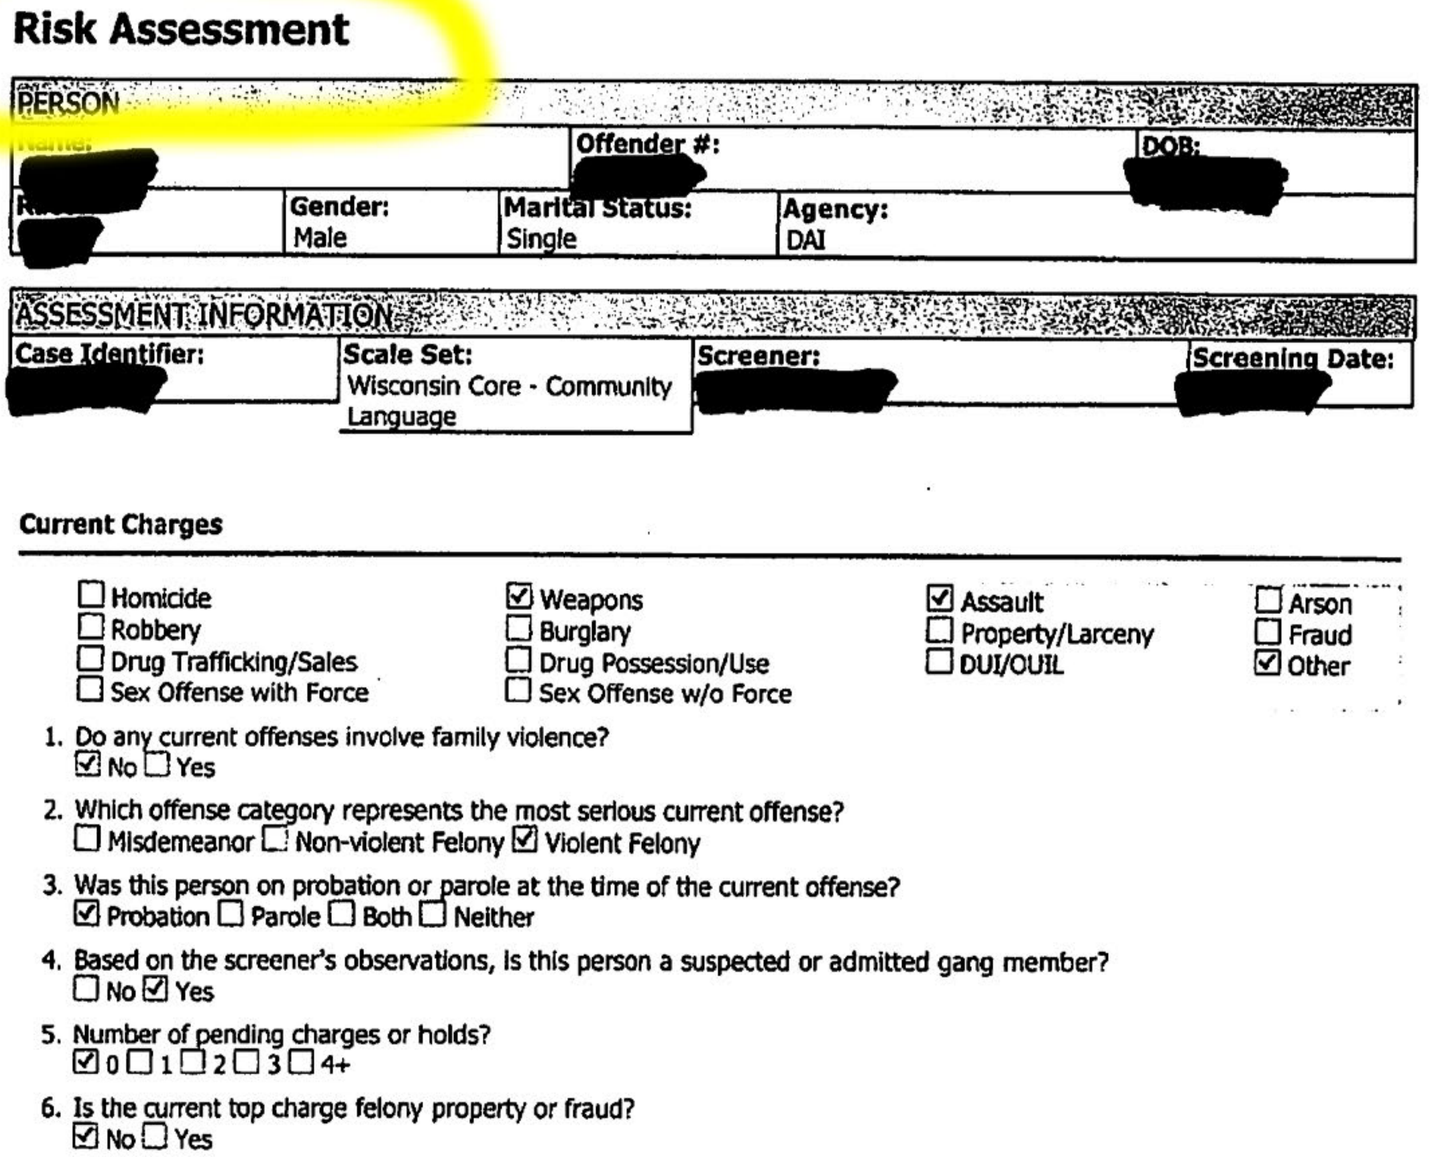

Risk Assessment

PERSON

[redacted]		Offender #: [redacted]	DOB: [redacted]
[redacted]	Gender: Male	Marital Status: Single	Agency: DAI

ASSESSMENT INFORMATION

Case Identifier: [redacted]	Scale Set: Wisconsin Core - Community Language	Screener: [redacted]	Screening Date: [redacted]

Current Charges

☐ Homicide	☑ Weapons	☑ Assault	☐ Arson
☐ Robbery	☐ Burglary	☐ Property/Larceny	☐ Fraud
☐ Drug Trafficking/Sales	☐ Drug Possession/Use	☐ DUI/OUIL	☑ Other
☐ Sex Offense with Force	☐ Sex Offense w/o Force		

1. Do any current offenses involve family violence?
   ☑ No ☐ Yes
2. Which offense category represents the most serious current offense?
   ☐ Misdemeanor ☐ Non-violent Felony ☑ Violent Felony
3. Was this person on probation or parole at the time of the current offense?
   ☑ Probation ☐ Parole ☐ Both ☐ Neither
4. Based on the screener's observations, is this person a suspected or admitted gang member?
   ☐ No ☑ Yes
5. Number of pending charges or holds?
   ☑ 0 ☐ 1 ☐ 2 ☐ 3 ☐ 4+
6. Is the current top charge felony property or fraud?
   ☑ No ☐ Yes

ALGORITHMIC FAIRNESS CASE #1 OF 2
INDIVIDUAL *versus* SOCIETY

RIGHTS OF THE ACCUSED
versus
RIGHTS EVERYONE ELSE

SCENARIOS NEXT TWO SLIDES:

Which is more fair to those accused?
Which is more fair for everyone else?
Which is more fair?

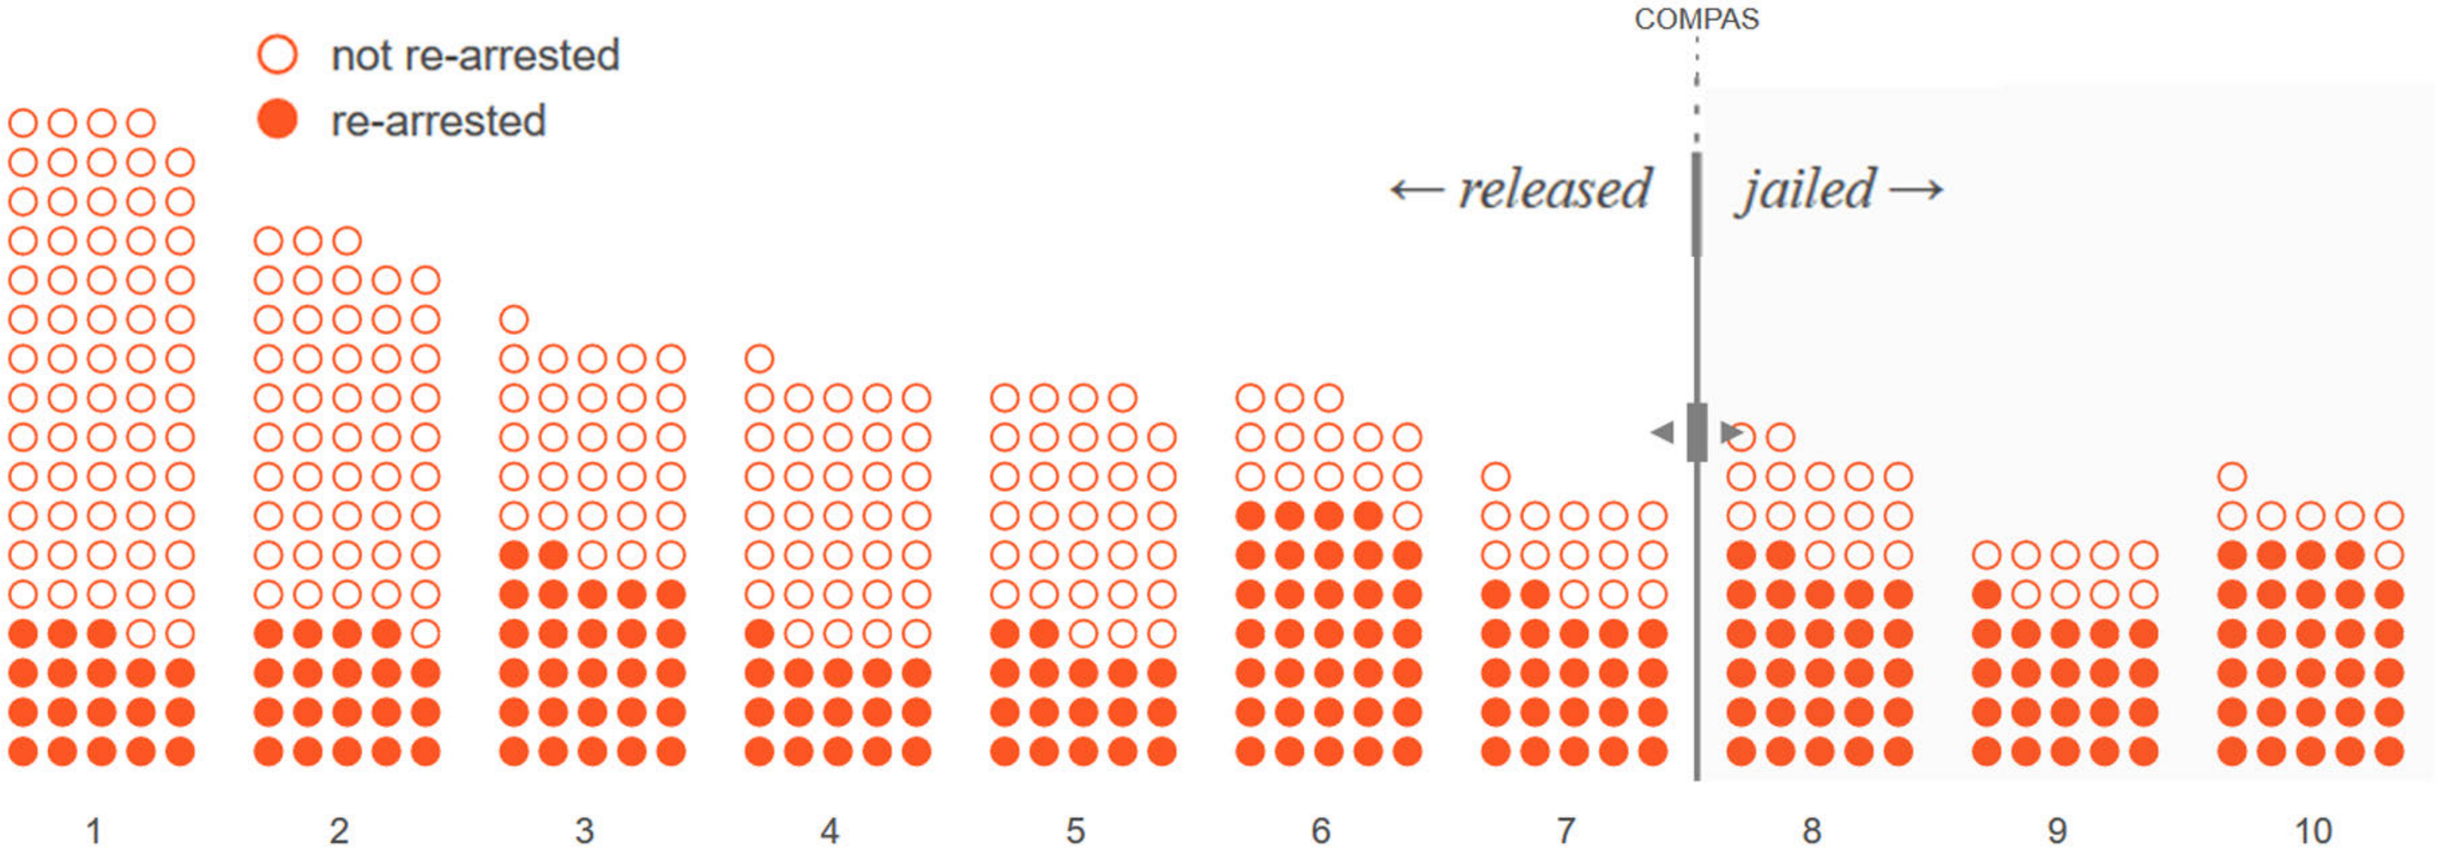

RELEASED BUT RE-ARRESTED

67%

Out of the 230 defendants re-arrested, 153 are rated "low risk."

NEEDLESSLY JAILED

11%

Out of the 270 defendants not re-arrested, 31 are rated "high risk."

RELEASED BUT RE-ARRESTED

35%

Out of the 230 defendants re-arrested, 80 are rated "low risk."

NEEDLESSLY JAILED

34%

Out of the 270 defendants not re-arrested, 91 are rated "high risk."

ALGORITHMIC FAIRNESS CASE #2 OF 2
Accused *versus* accused
COMPAS versus ProPublica

Access to fredom framed by recidivism, race, gender, age

COMPAS *versus* ProPublica
Background

ProPublica, a social activist organization, examined COMPAS statistics and found, correctly, that blacks who did not reoffend were proportionately more likely to be categorized as high risk than whites who did not reoffend.

COMPAS responded, correctly, that the recidivism score was uniform across races: blacks assigned 7 would reoffend at similar rates as whites assigned 7.

The key to the difference is a core statistic: in the data set, blacks overall offend at a higher rate than whites. Consequently, you can have uniformity across races in terms of calibration (how well the assigned number reflects recidivism risk), or parity (percentage of false needless imprisonments within each racial pool). But you can't have both.

Algorithmic fairness case #2 of 2

Statistical imbalances in crime rates occur across demographics: men offend more then women, the young more than the old.

The reasons for the divergences are in debate, and the specific numbers vary, but the theoretical ethical question is fixed.

Should people be judged – granted or denied freedom – as individuals? as members of a demographic? as a mix?

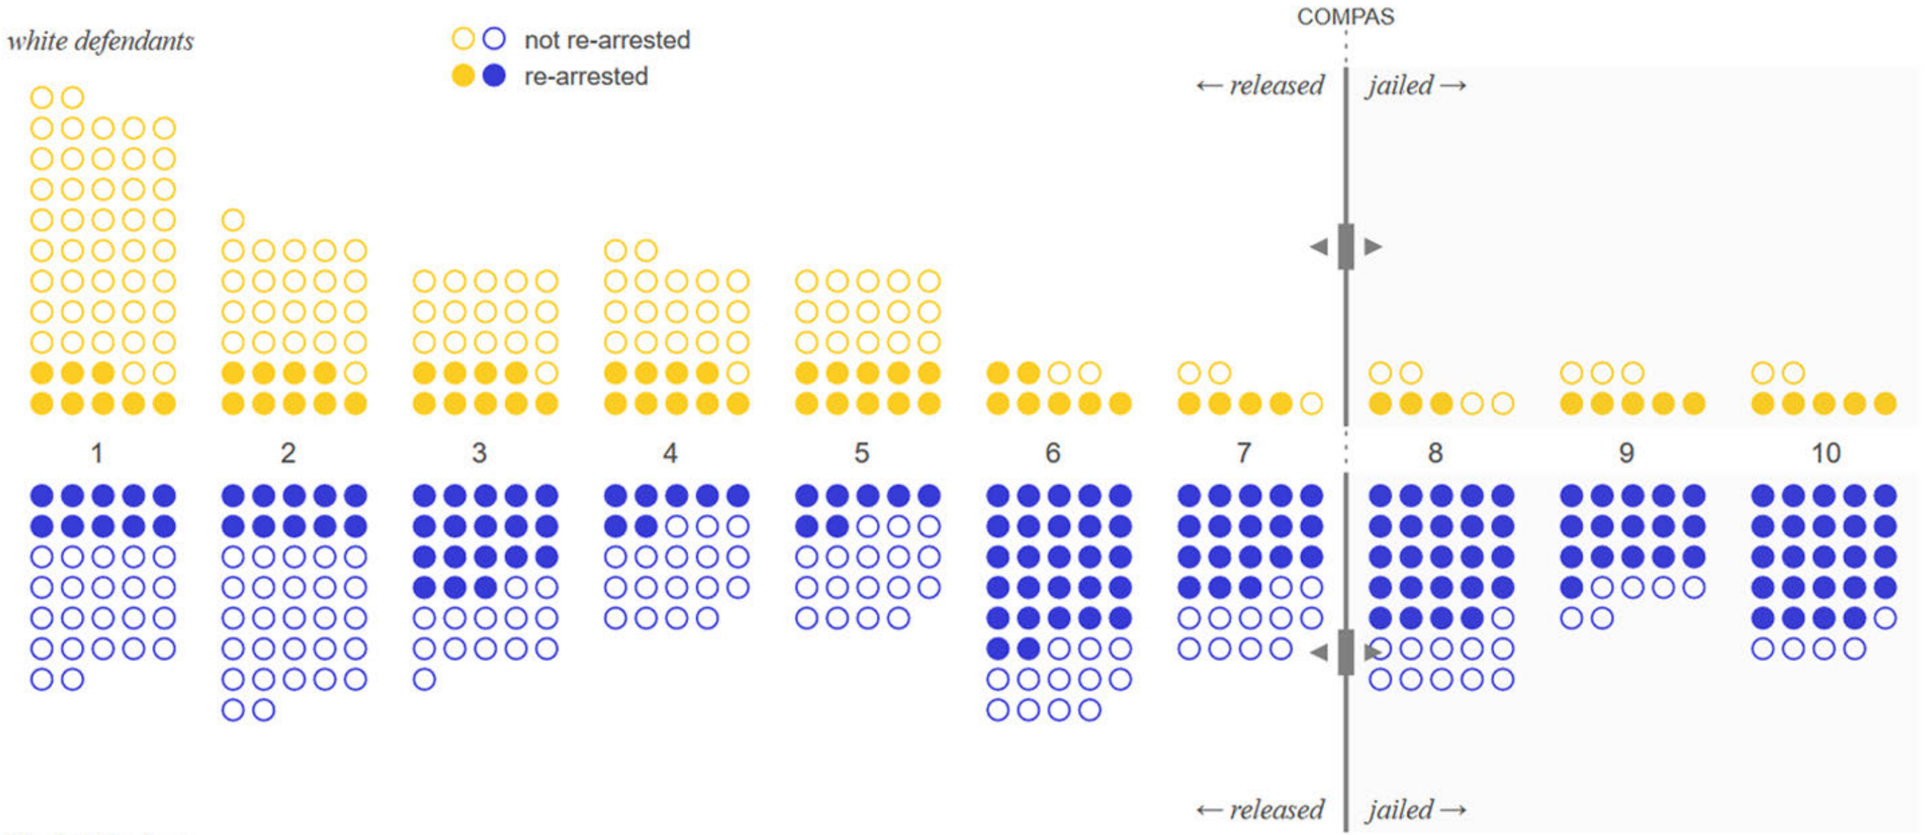

RELEASED BUT RE-ARRESTED

WHITE **81%**

Out of the 69 defendants re-arrested, 56 are rated "low risk."

BLACK **60%**

Out of the 161 defendants re-arrested, 97 are rated "low risk."

NEEDLESSLY JAILED

WHITE **7%**

Out of the 129 defendants not re-arrested, 9 are rated "high risk."

BLACK **16%**

Out of the 141 defendants not re-arrested, 22 are rated "high risk."

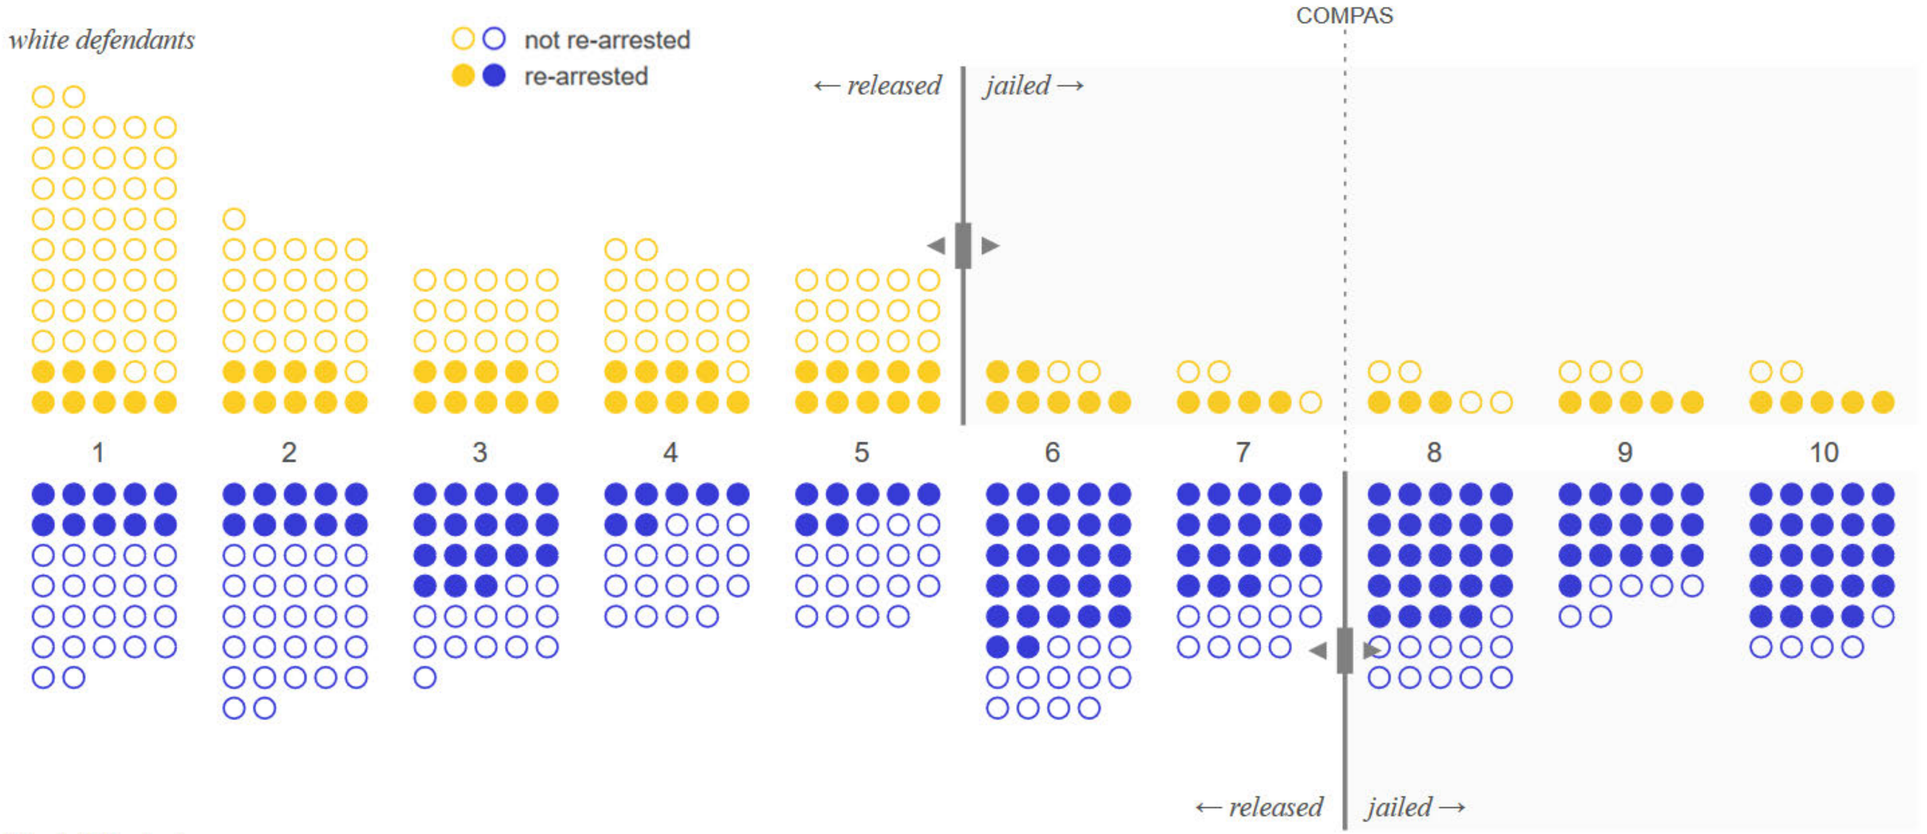

RELEASED BUT RE-ARRESTED

WHITE **65%**

Out of the 69 defendants re-arrested, 45 are rated "low risk."

BLACK **60%**

Out of the 161 defendants re-arrested, 97 are rated "low risk."

NEEDLESSLY JAILED

WHITE **11%**

Out of the 129 defendants not re-arrested, 14 are rated "high risk."

BLACK **16%**

Out of the 141 defendants not re-arrested, 22 are rated "high risk."

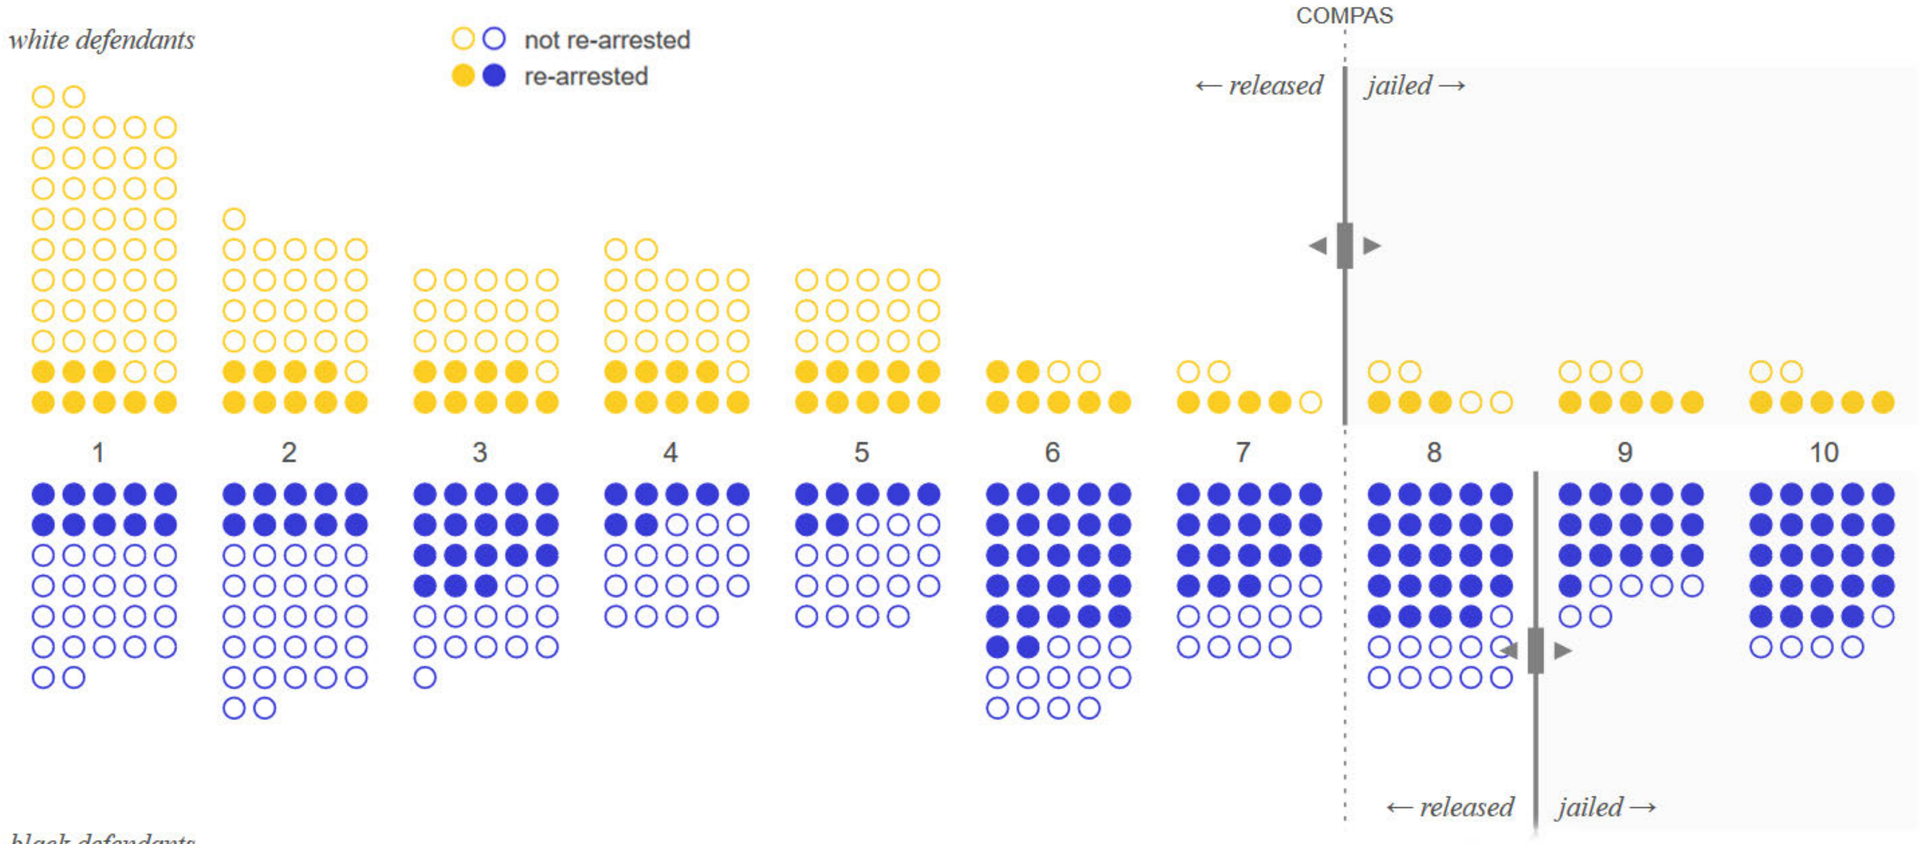

RELEASED BUT RE-ARRESTED

WHITE **81%**

Out of the 69 defendants re-arrested, 56 are rated "low risk."

BLACK **75%**

Out of the 161 defendants re-arrested, 121 are rated "low risk."

NEEDLESSLY JAILED

WHITE **7%**

Out of the 129 defendants not re-arrested, 9 are rated "high risk."

BLACK **8%**

Out of the 141 defendants not re-arrested, 11 are rated "high risk."

COMPETING LOGICS OF FAIRNESS

JUDGE/HUMAN

Decides cases *individually* as a verdict: Single best approximation, repeatedly.
Fairness applied to humans, gathered statistics later reflect what was done.

ALGORITHM/AI

Decides all cases simultaneously as a *range*: One decision applied to each subject, repeatedly.
Fairness applied to statistics, and humans later experience what was done.

IS ONE PREFERABLE TO THE OTHER?

James Brusseau
Philosophy Department
Pace University NYC

AI Ethics Site
Artificial Intelligence + Human Experience

Bernard Parker/Dylan Fugett
Josh Ritchie, for ProPublica
Public Domain

Graphic jailed/freed distributions
Can you make AI fairer than a judge? Play our courtroom algorithm game
MIT Technology Review, Oct 17, 2019
https://www.technologyreview.com/s/613508/ai-fairer-than-judge-criminal-risk-assessment-algorithm/