Bayesian Decision Making Lifts off with PyMC3

Thomas Wiecki, PhD

@twiecki

PyMC Labs: Bayesian consulting

Inventors of PyMC3, the leading platform for statistical data science

Decades of experience building Bayesian models

Team of:

- PhDs
- Mathematicians
- Neuroscientists
- Social scientists
- A former SpaceX rocket scientist

Adrian Seyboldt

Alexandre Andorra

Brandon Willard

Eric J. Ma

Luciano Paz

Maxim Kochurov

Oriol Abril Pla

Ravin Kumar

Thomas Wiecki

Used in industry

PyMC3

… and academia

Total citations

Cited by 1097

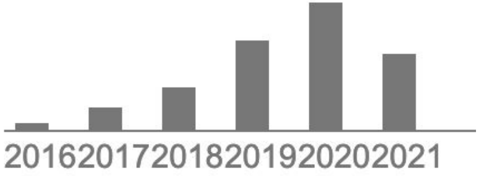

Probabilistic programming in Python using PyMC3
J Salvatier, TV Wiecki, C Fonnesbeck
PeerJ Computer Science 2, e55

[HTML] Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions
J Dehning, J Zierenberg, FP Spitzner, M Wibral… - …, 2020 - science.sciencemag.org

On the Fermi-GBM event 0.4 s after GW150914
J Greiner, JM Burgess, V Savchenko… - The Astrophysical …, 2016 - iopscience.iop.org

LUNA: quantifying and leveraging uncertainty in android malware analysis through Bayesian machine learning
M Backes, M Nauman - Security and Privacy (EuroS&P), 2017 …, 2017 - ieeexplore.ieee.org

[HTML] Dose-dependent regulation of alternative splicing by MBNL proteins reveals biomarkers for myotonic dystrophy
SD Wagner, AJ Struck, R Gupta, DR Farnsworth… - PLoS …, 2016 - journals.plos.org

[HTML] Confidence is higher in touch than in vision in cases of perceptual ambiguity
MT Fairhurst, E Travers, V Hayward, O Deroy - Scientific reports, 2018 - nature.com

Evaluation of Bayesian source estimation methods with Prairie Grass observations and Gaussian plume model: A comparison of likelihood functions and distance …
Y Wang, H Huang, L Huang, B Ristic - Atmospheric environment, 2017 -

Asymmetry in serial femtosecond crystallography data
A Sharma, L Johansson, E Dunevall… - … A: Foundations and …, 2017 - scripts.iucr.org

Seabirds enhance coral reef productivity and functioning in the absence of invasive rats
NAJ Graham, SK Wilson, P Carr, AS Hoey, S Jennings… - Nature, 2018 - nature.com

Limits on the number of spacetime dimensions from GW170817
K Pardo, M Fishbach, DE Holz, DN Spergel - arXiv preprint arXiv …, 2018 - arxiv.org

Blackbox ML vs Bayesian modeling

VS

- Pre-made, easy
- Can't customize
- One-size-fits-many
- Don't learn about ingredients
- More expensive (requires more data)

- Handmade, requires skill
- Can include dietary constraints (expert knowledge)
- Exactly to your taste
- Recipes can guide you
- Healthier ;-)

Insuring Rocket Launches


ひふみ

Data

Table 4: Selected current launch service providers

Vehicle	Launching state	Launch reliability 2008-18	Launch reliability %	Year of First Launch	Payload to LEO (kg)	Payload to GTO (kg)	Approximate cost per launch
Antares 230	USA	4/4	100%	2016	7,000	2,700	$271.5m
Atlas V 401	USA	32/32	100%	2002	9,797	4,750	$132m - $164m
Atlas V 541	USA	6/6	100%	2011	17,410	8,290	$243
Delta IV Medium+ (5,4)	USA	7/7	100%	2009	14,140	6,337	$137m
Falcon 9 Upgrade (v1.2)	USA	47/47	100%	2015	22,800	8,300	$62m
Falcon Heavy	USA	1/1	100%	2018	63,800	26,700	$90m
Proton M Briz M	Russia	70/76	92%	2001	23,000	6,920	$105m
Rokot	Russia	20/21	95%	1994	2,140		$30m
Soyuz 2-1A	Russia	26/28	93%	2004	7,400	1,500	$46m
Soyuz 2-1B	Russia	25/27	93%	2006	8,250	1,800	$46m
Soyuz-FG	Russia	44/45	98%	2001	7,200		
Long March 2C	China	24/25	96%	1975	3,850	1,250	
Long March 2D	China	33/34	97%	1992	4,000		
Long March 3B	China	21/22	95%	1996	13,600	5,100	
Long March 3BE	China	21/22	95%	2007		5,500	
Long March 4B	China	20/21	95%	1999	2,230		
Long March 4C	China	22/23	96%	2006	2,950	1,500	
Ariane V ECA	Europe	55/56	98%	1996	21,000	10,000	$137m
Ariane V ES/ATV	Europe	8/8	100%	2008	20,000	8,000	$137m
Soyuz ST-A	Europe	6/6	100%	2011	4,340	2,760	$73m - $78m
Soyuz ST-B	Europe	13/14	93%	2011	4,900	3,150	$73m - $78m
Vega	Europe	12/12	100%	2012	1,500		$46m
GSLV Mk II	India	4/5	80%	2007	5,000	2,500	$40m
GSLV Mk III	India	2/2	100%	2017	3,000	4,000	$60m
PSLV XL	India	18/19	95%	2008	1,700	1,425	$22m
H-IIA 202	Japan	23/23	100%	2001	3,300	4,000	$82m
GSLV Mk II	India	4/5	80%	2007	7,000	2,700	$40m

Source: Space Foundation (2018), The Space Report 2018 and London Economics analysis

Problem setting

- Fixed budget we want to allocate
- How to distribute?
- Those with 100% reliability seem like the safest bet
- Antares 230 and Atlas V 401 both have 100% reliability, so they are same, right?
- What's missing: **uncertainty quantification**

Table 4: Selected current launch service providers

Vehicle	Launching state	Launch reliability 2008-18	Launch reliability %	Year of First Launch	Payload to LEO (kg)	Payload to GTO (kg)	Approximate cost per launch
Antares 230	USA	4/4	100%	2016	7,000	2,700	$271.5m
Atlas V 401	USA	32/32	100%	2002	9,797	4,750	$132m - $164m
Atlas V 541	USA	6/6	100%	2011	17,410	8,290	$243
Delta IV Medium+ (5,4)	USA	7/7	100%	2009	14,140	6,337	$137m
Falcon 9 Upgrade (v1.2)	USA	47/47	100%	2015	22,800	8,300	$62m
Falcon Heavy	USA	1/1	100%	2018	63,800	26,700	$90m
Proton M Briz M	Russia	70/76	92%	2001	23,000	6,920	$105m
Rokot	Russia	20/21	95%	1994	2,140		$30m
Soyuz 2-1A	Russia	26/28	93%	2004	7,400	1,500	$46m
Soyuz 2-1B	Russia	25/27	93%	2006	8,250	1,800	$46m
Soyuz-FG	Russia	44/45	98%	2001	7,200		
Long March 2C	China	24/25	96%	1975	3,850	1,250	
Long March 2D	China	33/34	97%	1992	4,000		
Long March 3B	China	21/22	95%	1996	13,600	5,100	
Long March 3BE	China	21/22	95%	2007		5,500	
Long March 4B	China	20/21	95%	1999	2,230		
Long March 4C	China	22/23	96%	2006	2,950	1,500	
Ariane V ECA	Europe	55/56	98%	1996	21,000	10,000	$137m
Ariane V ES/ATV	Europe	8/8	100%	2008	20,000	8,000	$137m
Soyuz ST-A	Europe	6/6	100%	2011	4,340	2,760	$73m - $78m
Soyuz ST-B	Europe	13/14	93%	2011	4,900	3,150	$73m - $78m
Vega	Europe	12/12	100%	2012	1,500		$46m
GSLV Mk II	India	4/5	80%	2007	5,000	2,500	$40m
GSLV Mk III	India	2/2	100%	2017	3,000	4,000	$60m
PSLV XL	India	18/19	95%	2008	1,700	1,425	$22m
H-IIA 202	Japan	23/23	100%	2001	3,300	4,000	$82m
GSLV Mk II	India	4/5	80%	2007	7,000	2,700	$40m

Source: Space Foundation (2018), The Space Report 2018 and London Economics analysis

Quantifying uncertainty with Bayesian modeling

Instead of specifying the most likely value (e.g. 100%), we **assign beliefs to every possible state** (0% to 100%) using a probability distribution.

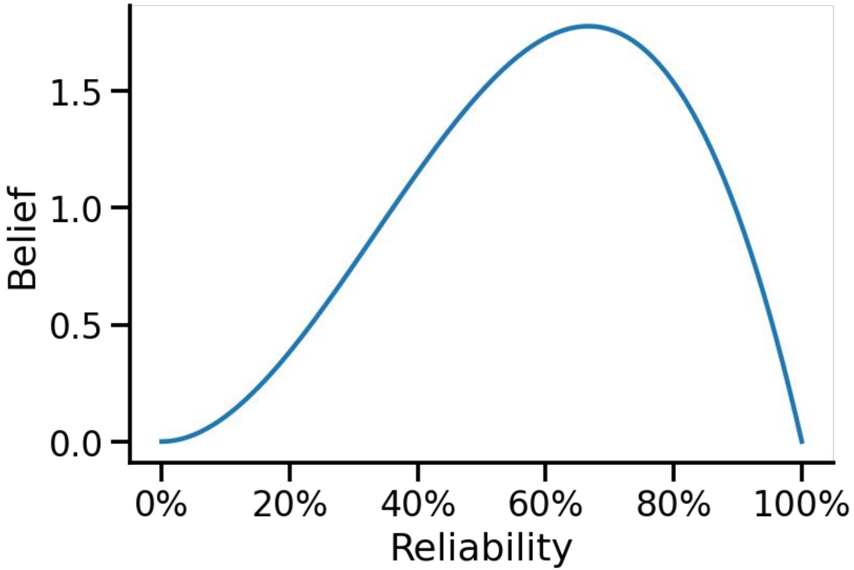

Priors

Before we look at any data, we first specify our beliefs in all possible states using a **prior distribution**.

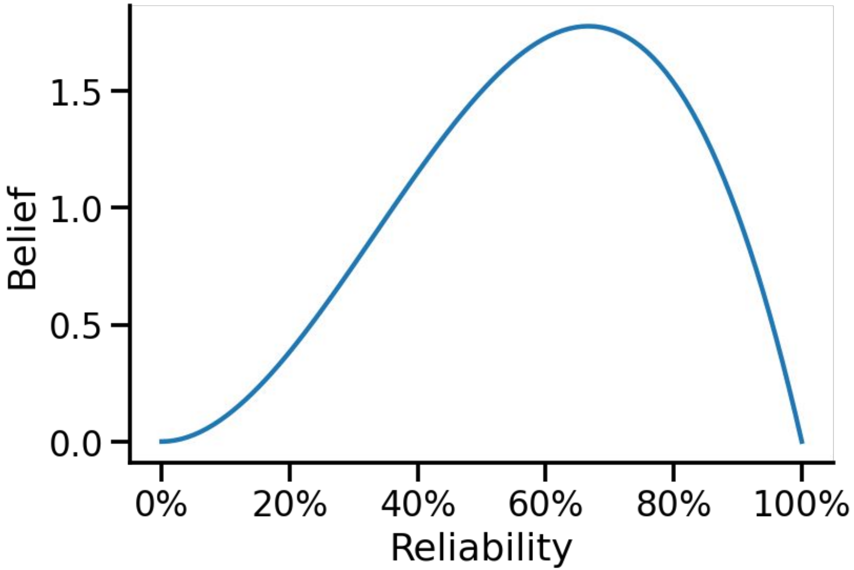

Posterior distribution

When we see data, we **update our beliefs** about the possible states. The more data we observe, the more concentrated our beliefs will be.

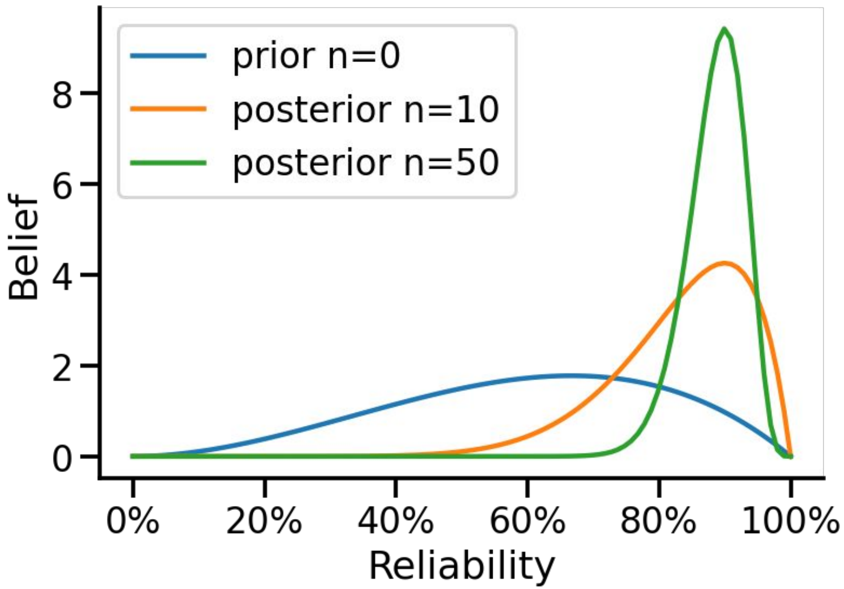

Modeling our data

- Our data is successes out of total trials → binomial distribution
- This distribution

Proton M Briz M	Russia	70/76
Rokot	Russia	20/21
Soyuz 2-1A	Russia	26/28
Soyuz 2-1B	Russia	25/27
Soyuz-FG	Russia	44/45

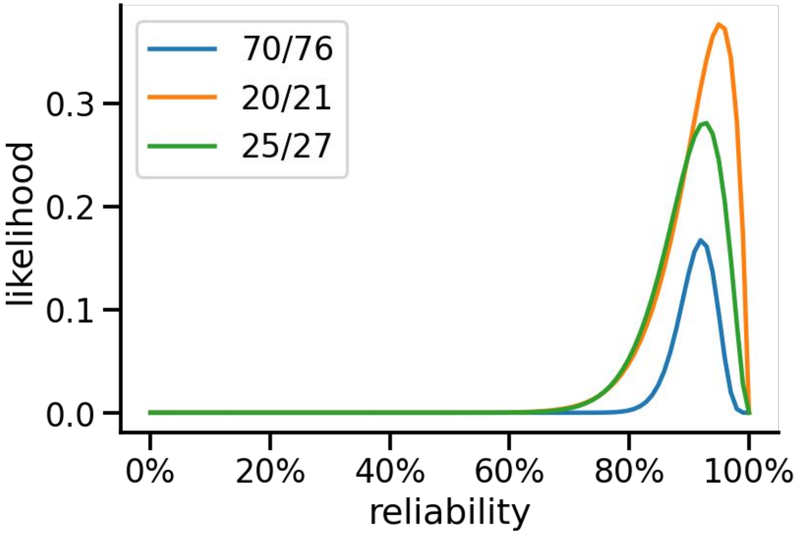

A Tale of Two Spaces

Parameter space

What we want to infer

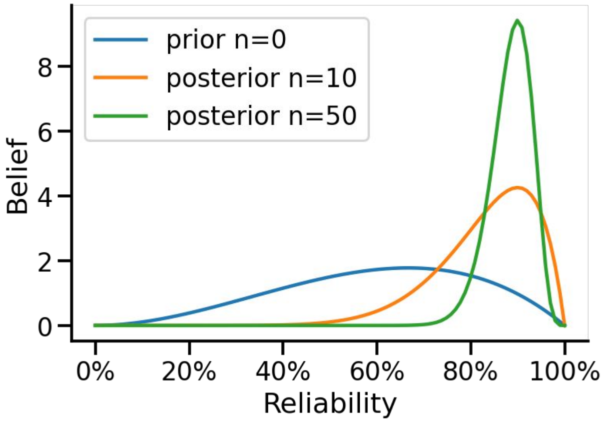

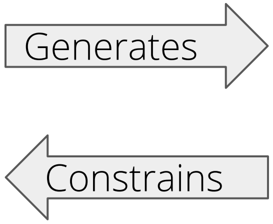

Data space

What we observe

Proton M Briz M	Russia	70/76
Rokot	Russia	20/21
Soyuz 2-1A	Russia	26/28
Soyuz 2-1B	Russia	25/27
Soyuz-FG	Russia	44/45

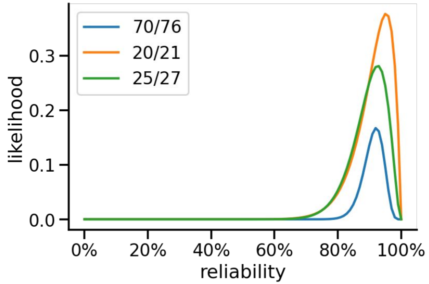

Getting data into Python

vehicle	country	successes	total	percentage	first_year	leo	gto	cost	prob	payoff
Antares 230	USA	4	4	100	2016	7000	2700.0	271.5	0.744	162.9
Atlas V 401	USA	32	32	100	2002	9797	4750.0	148.0	0.870	88.8
Atlas V 541	USA	6	6	100	2011	17410	8290.0	243.0	0.758	145.8
Delta IV Medium+ (5.4)	USA	7	7	100	2009	14140	6337.0	137.0	0.750	82.2
Falcon 9 Upgrade (v1.2)	USA	47	47	100	2015	22800	8300.0	62.0	0.891	37.2
Falcon Heavy	USA	1	1	100	2018	63800	26700.0	90.0	0.704	54.0
Proton M Briz M	Russia	70	76	92	2001	23000	6920.0	105.0	0.845	63.0
Rokot	Russia	20	21	95	1994	2140	NaN	30.0	0.798	18.0
Soyuz 2-1A	Russia	26	28	93	2004	7400	1500.0	46.0	0.802	27.6
Soyuz 2-1B	Russia	25	27	93	2006	8250	1800.0	46.0	0.800	27.6
Ariane V ECA	Europe	55	56	98	1996	21000	10000.0	137.0	0.885	82.2
Ariane V ES/ATV	Europe	8	8	100	2008	20000	8000.0	137.0	0.770	82.2
Soyuz ST-A	Europe	6	6	100	2011	4340	2760.0	75.5	0.756	45.3
Soyuz ST-B	Europe	13	14	93	2011	4900	3150.0	75.5	0.776	45.3
Vega	Europe	12	12	100	2012	1500	NaN	46.0	0.801	27.6
GSLV Mk II	India	4	5	80	2007	5000	2500.0	40.0	0.698	24.0
GSLV Mk III	India	2	2	100	2017	3000	4000.0	60.0	0.726	36.0
PSLV XL	India	18	19	95	2008	1700	1425.0	22.0	0.797	13.2
H-IIA 202	Japan	23	23	100	2001	3300	4000.0	82.0	0.843	49.2

This is the intuition behind Bayesian statistics

1. Start with some belief about possible states of the world (Prior)
2. Combine with an intuition of how the world works (Model and Likelihood)
3. Update your beliefs as data comes in - some beliefs might not be plausible anymore (Posterior)

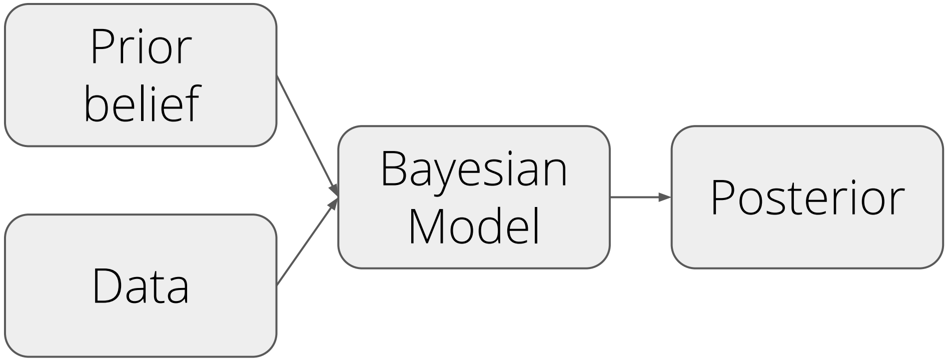

Here’s the model in PyMC3

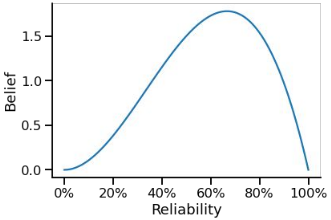

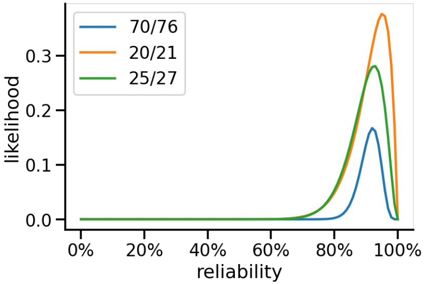

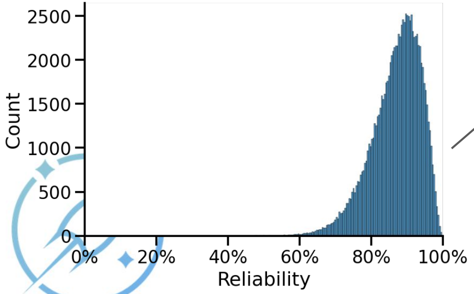


```python
import pymc3 as pm

with pm.Model() as model:
    # Define priors
    p = pm.Beta("p", alpha=6, beta=1,
                     shape=len(df))

    # Define likelihood
    obs = pm.Binomial("obs", p=p,
                           n=df.total.values,
                           observed=df.successes.values)

    # Run MCMC inference
    posterior = pm.sample()
```

total
4
32
6

successes
4
32
6

Results

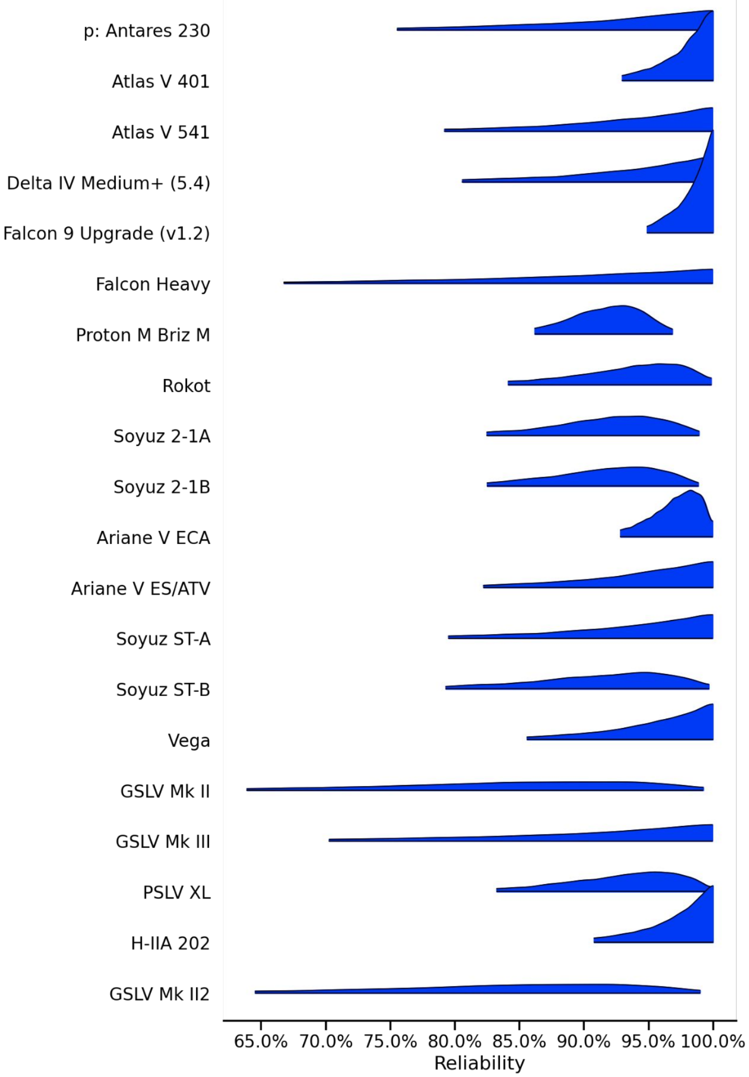

vehicle	successes	total
Antares 230	4	4
Atlas V 401	32	32
Atlas V 541	6	6
Delta IV Medium+ (5.4)	7	7
Falcon 9 Upgrade (v1.2)	47	47
Falcon Heavy	1	1
Proton M Briz M	70	76
Rokot	20	21
Soyuz 2-1A	26	28
Soyuz 2-1B	25	27
Ariane V ECA	55	56
Ariane V ES/ATV	8	8
Soyuz ST-A	6	6
Soyuz ST-B	13	14
Vega	12	12
GSLV Mk II	4	5
GSLV Mk III	2	2
PSLV XL	18	19
H-IIA 202	23	23
GSLV Mk II2	4	5

Models can be much more accurate

- Now that we have a simple model in place, it's a good idea to improve it.
- **PyMC3 makes this easy** as we just have to **extend the code**, no new derivations of estimators necessary.
- One example that could be useful here: use a **hierarchical model**
- This would estimate a group distribution for each country and exploit the similarities

Model ignoring similarities

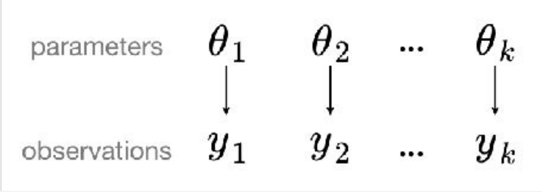

Hierarchical model with group distribution per country

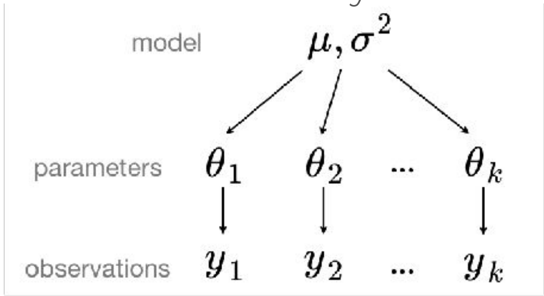

Let's instead go into a different direction.

Have we actually solved anything?

- Instead of just a single number, we now have **posterior distributions quantifying our uncertainty,** that's kinda cool.
- Most data science would just call it a day.
- However, for data science to have an impact on the bottom line: Rather than provide plots that may inform a decision, help **make a decision.**
- **Bayesian Decision Making** provides an elegant framework for this.

Decision Time

How do we make the decision that maximizes profit given our model estimates?

Step 1: Generate multiple plausible scenarios

Turn model parameters into scenarios according to their plausibility based on the data we have seen and the model.

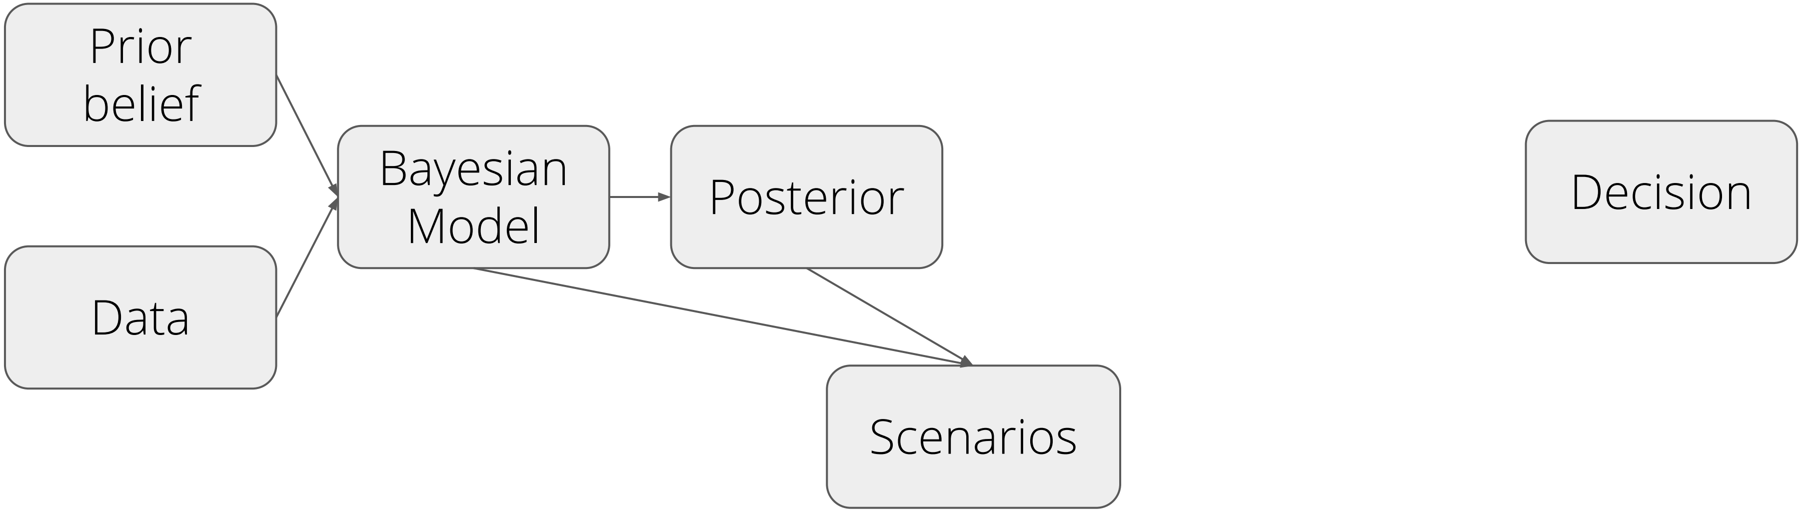

vehicle	Antares 230	Atlas V 401	Atlas V 541	Delta IV Medium+ (5.4)	Falcon 9 Upgrade (v1.2)	Falcon Heavy	Proton M Briz M	Rokot	Soyuz 2-1A	Soyuz 2-1B	Ariane V ECA	Ariane V ES/ATV
simulated launch												
0	1	0	1	1	1	1	1	1	0	1	1	1
1	1	1	0	1	1	1	1	1	1	1	1	1
2	1	1	1	1	1	1	1	1	0	1	1	1
3	1	1	1	1	1	0	1	1	1	1	1	1
4	0	1	1	1	1	1	1	1	1	1	1	0
...	...	...	...	...	...	...	...	...	...	...	...	...
995	1	1	1	1	1	1	1	1	1	1	1	1
996	1	1	1	1	1	1	1	1	1	1	1	1
997	0	1	1	1	1	1	1	1	1	1	1	1
998	1	1	1	1	0	1	0	1	1	1	1	1
999	0	1	1	1	1	1	1	0	1	1	1	1

Assign outcomes to scenarios

- Very simple assumptions:
  - If the rocket explodes, we lose the total cost of sending it to space (we have this from the able).
  - If the rocket lifts off, we get paid 60% of that total cost.
- We can easily make this more complicated, this is just for demonstration purposes.

vehicle	Antares 230	Atlas V 401	Atlas V 541	Delta IV Medium+ (5.4)	Falcon 9 Upgrade (v1.2)	Falcon Heavy	Proton M Briz M	Rokot	Soyuz 2-1A	Soyuz 2-1B	Ariane V ECA	Ariane V ES/ATV
simulated launch												
0	162.9	-148.0	145.8	82.2	37.2	54.0	63.0	18.0	-46.0	27.6	82.2	82.2
1	162.9	88.8	-243.0	82.2	37.2	54.0	63.0	18.0	27.6	27.6	82.2	82.2
2	162.9	88.8	145.8	82.2	37.2	54.0	63.0	18.0	-46.0	27.6	82.2	82.2
3	162.9	88.8	145.8	82.2	37.2	-90.0	63.0	18.0	27.6	27.6	82.2	82.2
4	-271.5	88.8	145.8	82.2	37.2	54.0	63.0	18.0	27.6	27.6	82.2	-137.0
...	...	...	...	...	...	...	...	...	...	...	...	...
995	162.9	88.8	145.8	82.2	37.2	54.0	63.0	18.0	27.6	27.6	82.2	82.2
996	162.9	88.8	145.8	82.2	37.2	54.0	63.0	18.0	27.6	27.6	82.2	82.2
997	-271.5	88.8	145.8	82.2	37.2	54.0	63.0	18.0	27.6	27.6	82.2	82.2
998	162.9	88.8	145.8	82.2	-62.0	54.0	-105.0	18.0	27.6	27.6	82.2	82.2
999	-271.5	88.8	145.8	82.2	37.2	54.0	63.0	-30.0	27.6	27.6	82.2	82.2

What's the profit *taking uncertainty into account*?

In order to find the best decision we need to define what *best* means by specifying an objective function.

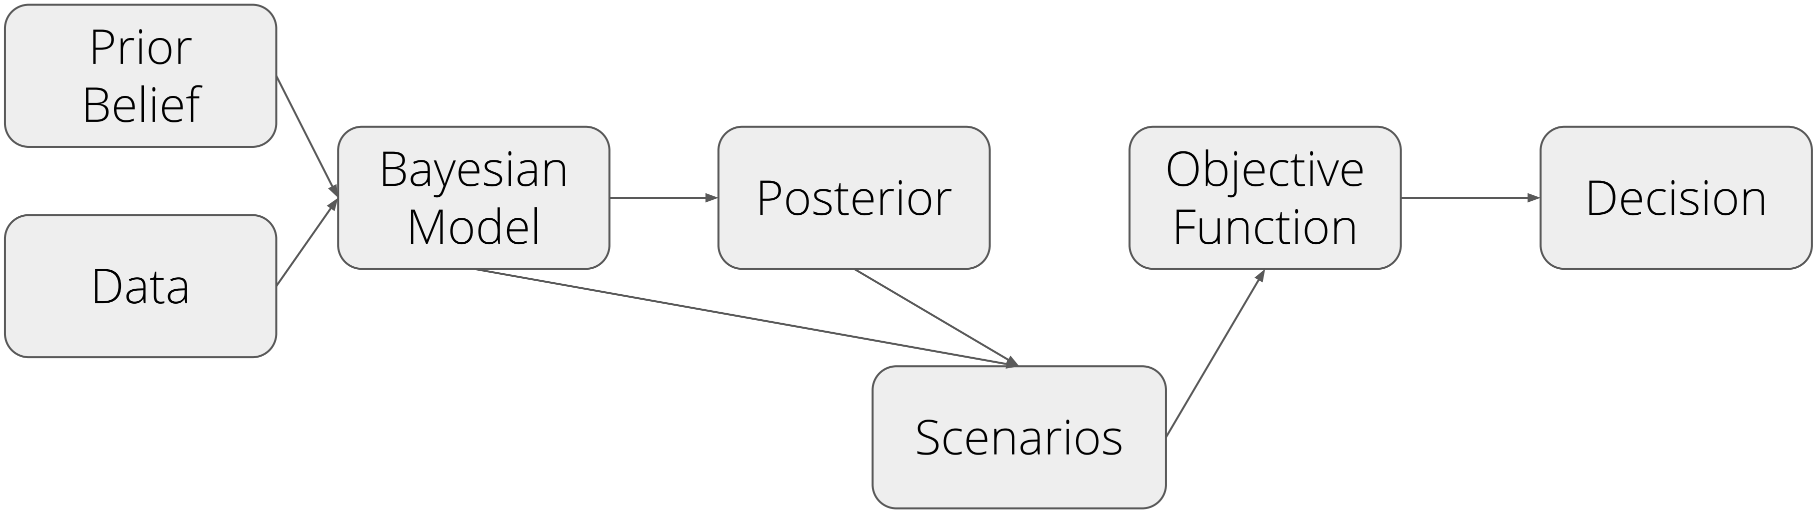

How should we allocate our budget?

Find order amount which **maximizes** profit across all simulated rocket launches while taking **constraints** (budget and max order size) into account.

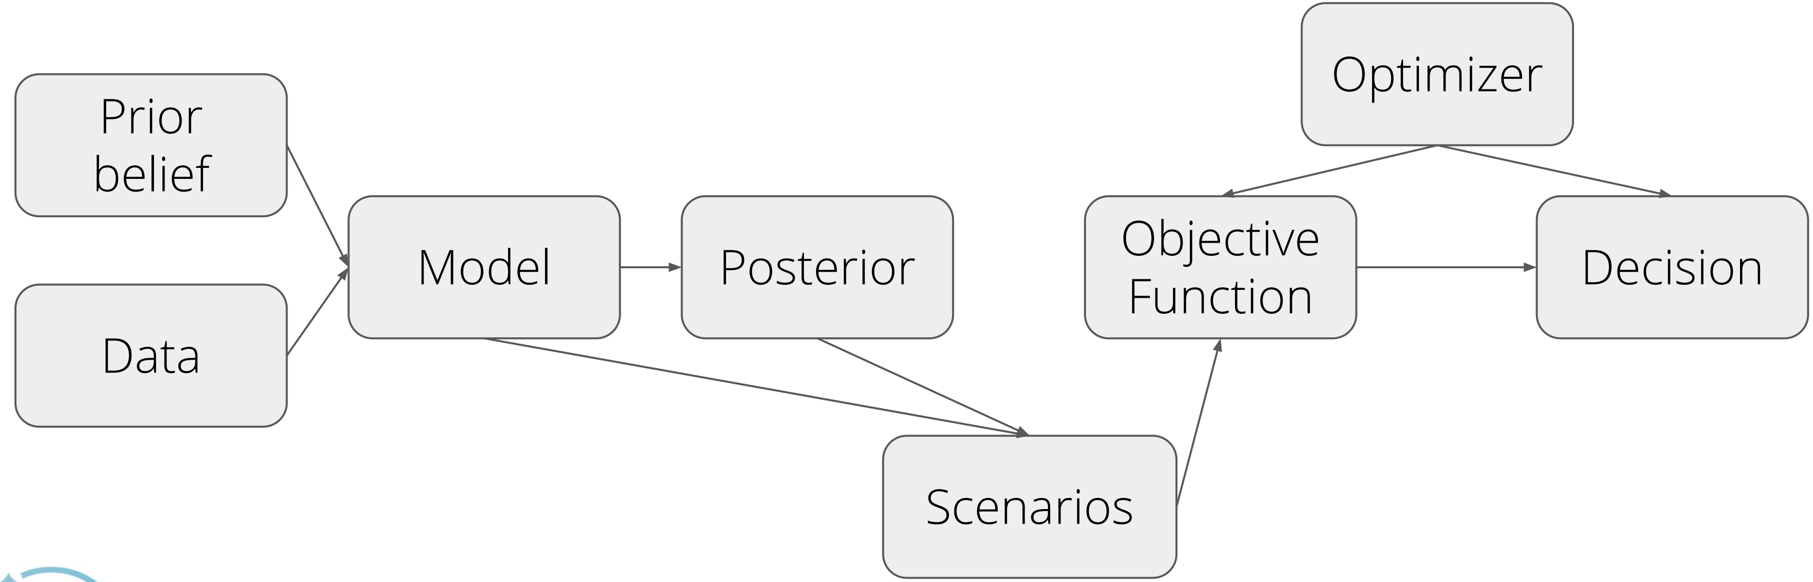

Pseudo-code (simplistic)

```
def compute_expected_profit(alloc): # e.g.: [.3, .2, .5]
    payoff = alloc * df_outcomes
    expected_payoff = mean(sum(payoff))
    return expected_payoff

optimal_alloc = optimizer.maximize(compute_expected_profit)
```

Optimal allocation across all scenarios

Posteriors (parameter space)

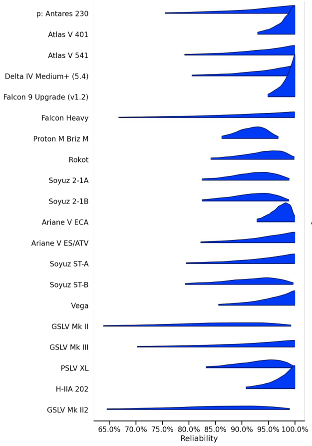

Outcomes (data space)

vehicle / simulated launch	Antares 230	Atlas V 401	Atlas V 541
0	162.9	-148.0	145.8
1	162.9	88.8	-243.0
2	162.9	88.8	145.8
3	162.9	88.8	145.8
4	-271.5	88.8	145.8
...	...	...	...
995	162.9	88.8	145.8
996	162.9	88.8	145.8
997	-271.5	88.8	145.8
998	162.9	88.8	145.8
999	-271.5	88.8	145.8

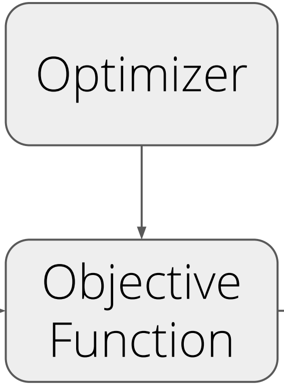

Decision

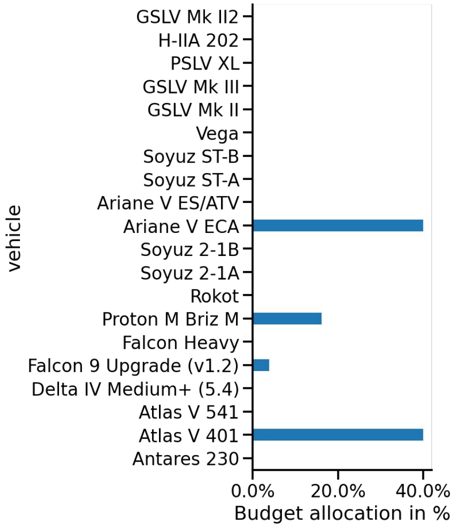

So how much profit are we expecting?

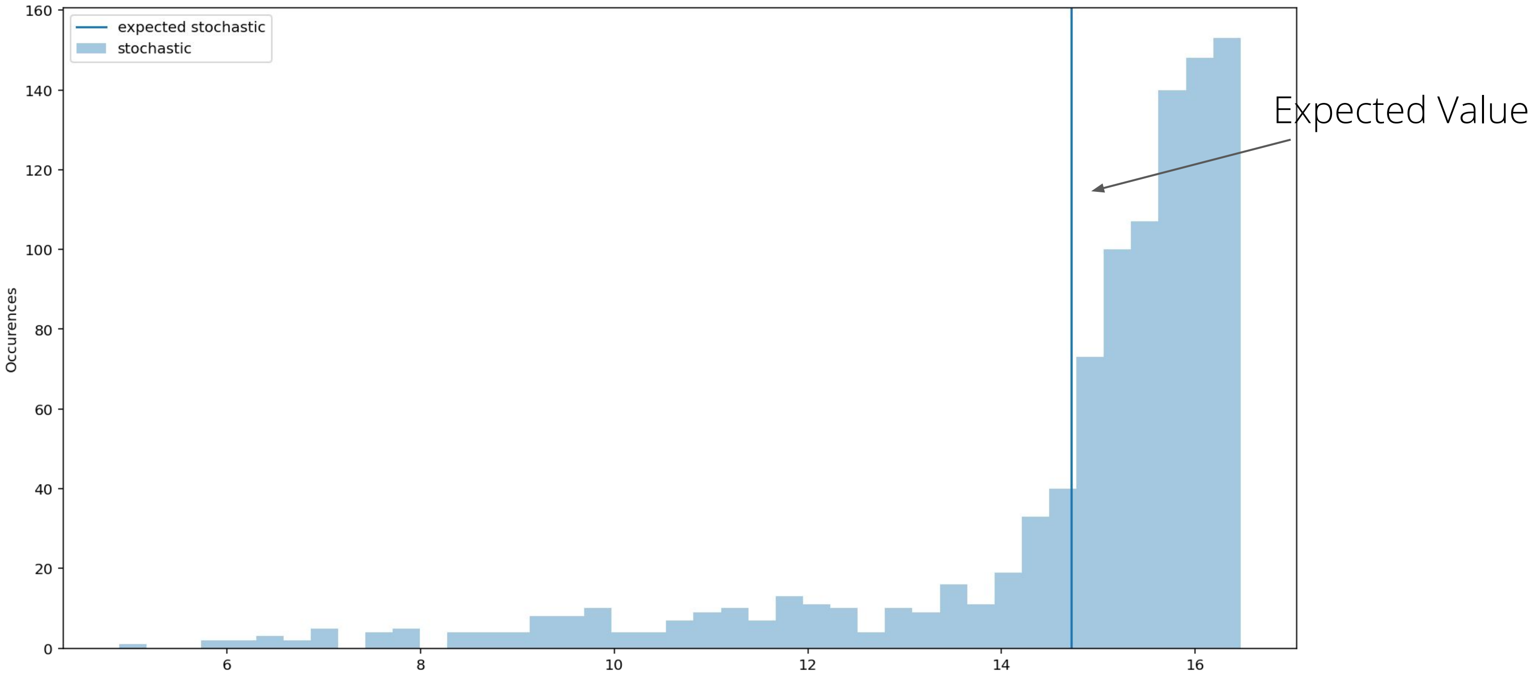

As we can't know when a rocket will crash, the outcome of our optimized decision will also be stochastic.

And what would be the outcome if we just used point estimates?

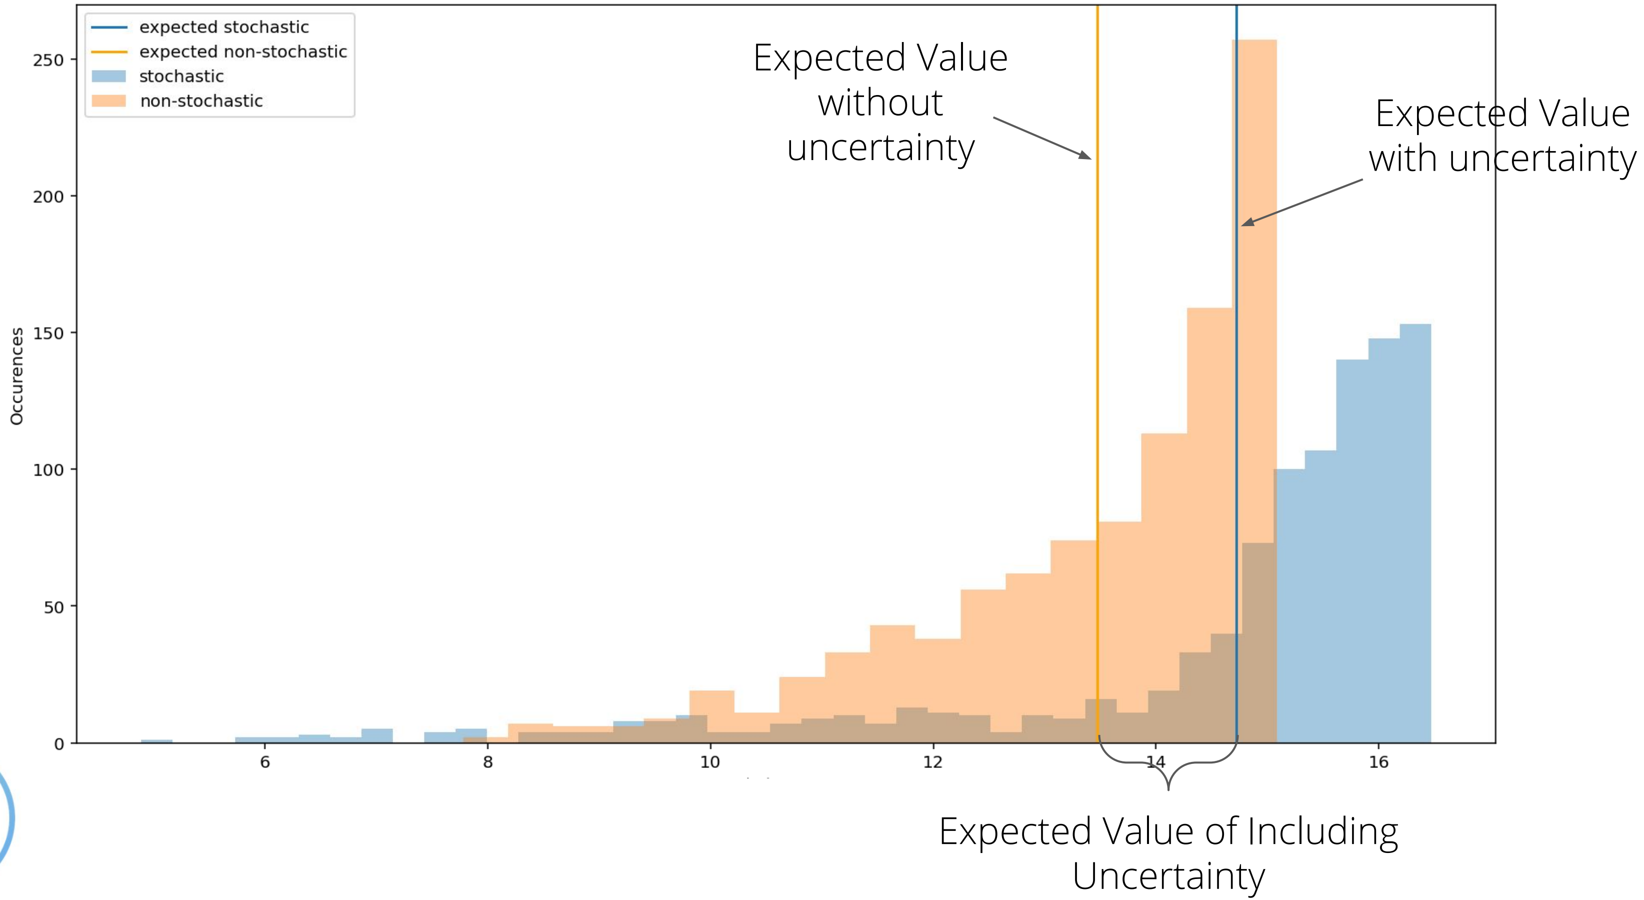

Benefits of Bayesian Model

- More robust as distributions are leveraged rather than point-estimates
    - The average doesn't tell you a whole lot about all the possibilities
- Different "track records" are automatically handled
    - Short but great track-record: high uncertainty → many potentially bad outcomes → low weight
- Framework: Model and objective can be improved to include all kinds of structure:
    - Hierarchical information about country/manufacturer
    - Risk-aversion
    - Payload
    - Estimate optimal insurance premia

Bayesian Insurance Data Science

- Insurance statistics is stuck in the past.
- The room for innovation is huge, Bayesian modeling perfect tool.
- → The possibility for disruption is huge. Be part of the future.
- We are looking for partners to create that future.

Resources

- PyMC3: www.pymc.io
- PyMC Labs: www.pymc-labs.io
- Blog post on Bayesian Decision Making: https://twiecki.io/blog/2019/01/14/supply_chain/