Learning and equilibrium in misspecified models

by Filippo Massari & Jonathan Newton

No statistical model is "true" or "false", "right" or "wrong"; the models just have varying performance, which can be assessed.

Jorma Rissanen - Information and complexity in statistical modeling (2007).

The human mind is programmed for survival, not for truth.

John Gray - Seven types of atheism (2018).


Canon

Consider a biased coin that has a true probability of heads of 0.7

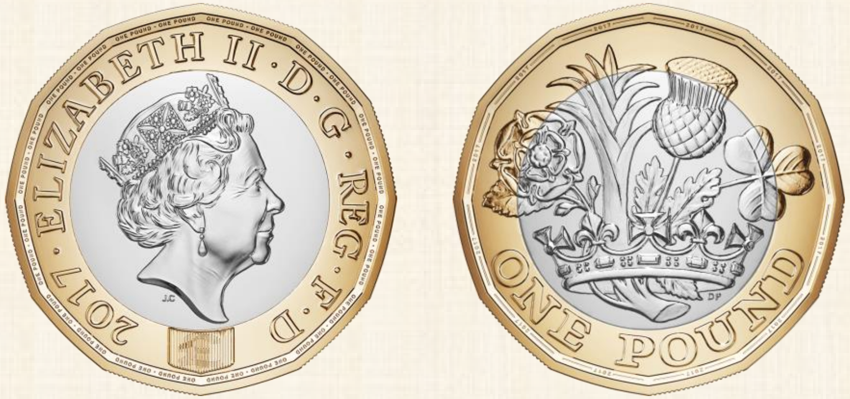

0.7 0.3

However Zhang only considers the following two possibilities, neither of which corresponds to the truth.

Her model is **misspecified**.

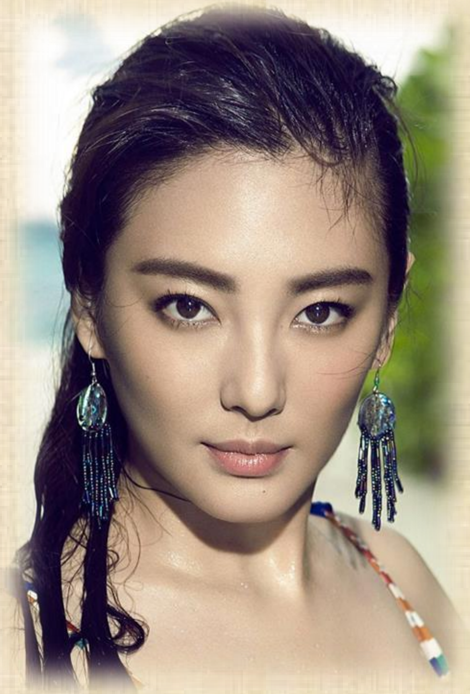

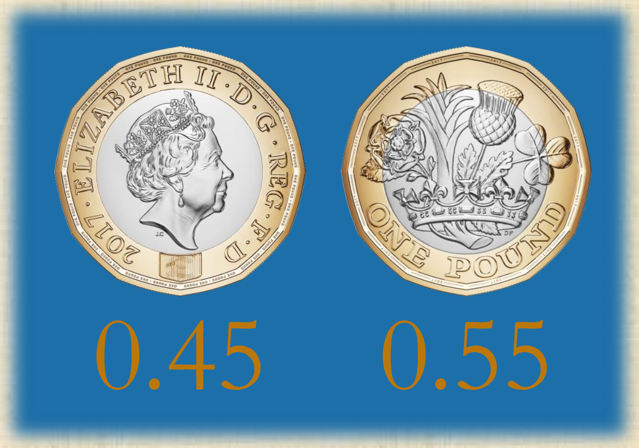

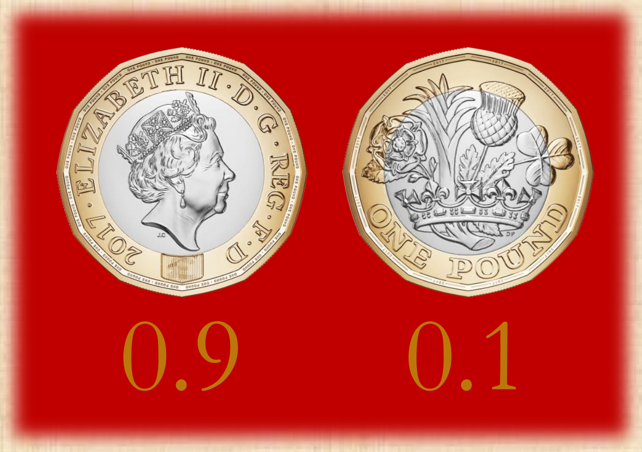

She observes a sequence of tosses of the coin and can bet **heads** or **tails** each time, earning a dollar every time she is correct.

Recall that the true probability of heads is 0.7

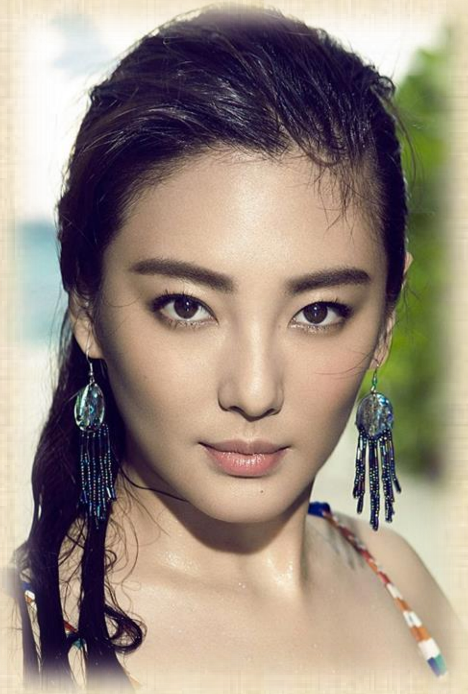

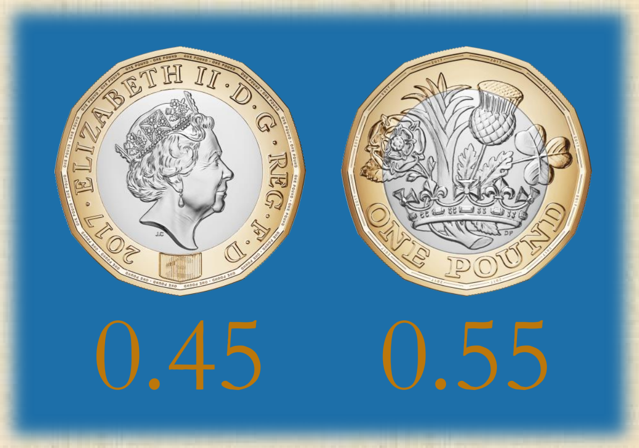

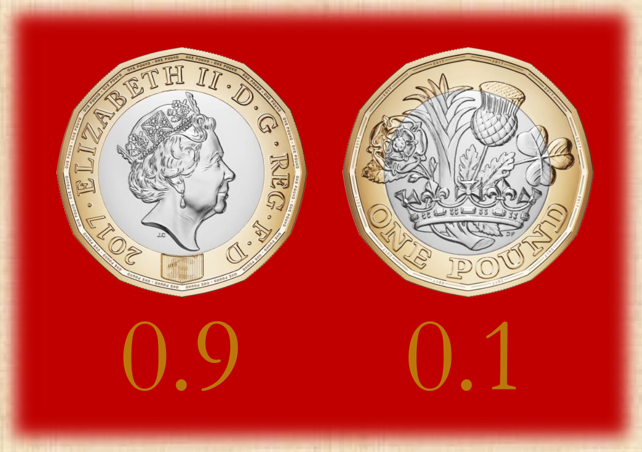

Zhang is a Bayesian learner.

If she places positive initial probability on the **blue** and the **red** models, then over time she places almost all weight on the **blue** model.

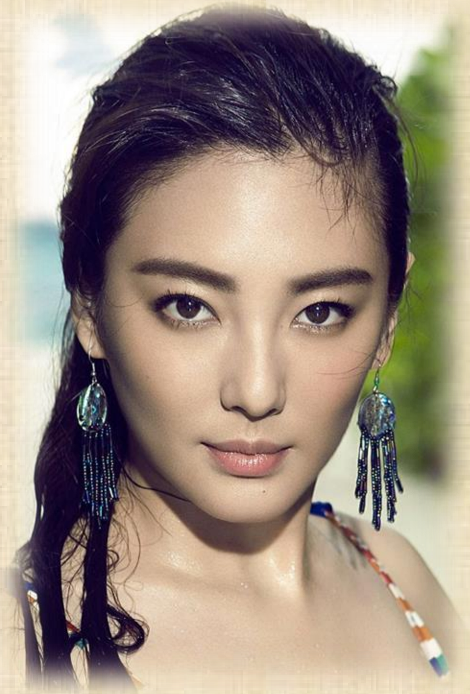

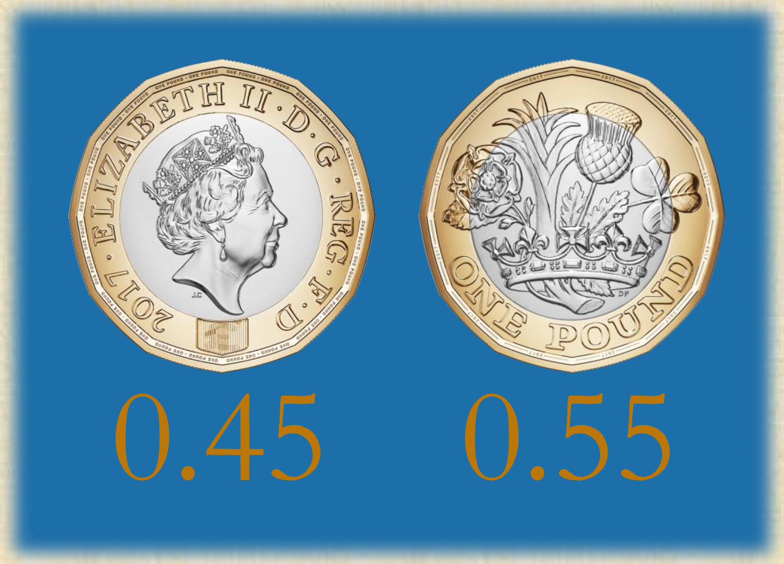

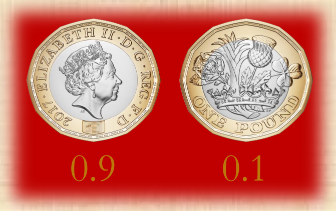

Hence she will bet **tails** every period.

As the true probability of **heads** is 0.7, she will earn an average payoff over time of 0.3 dollars per period.

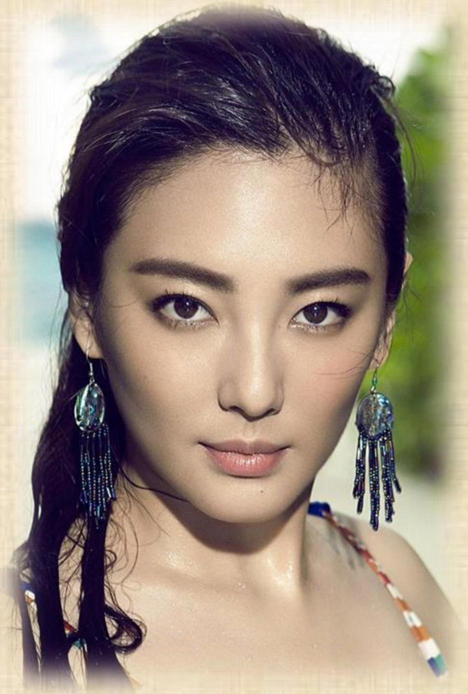

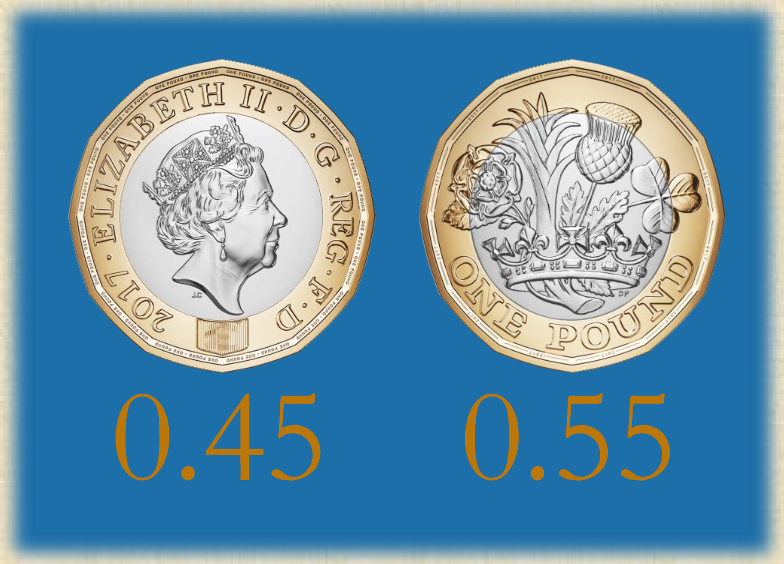

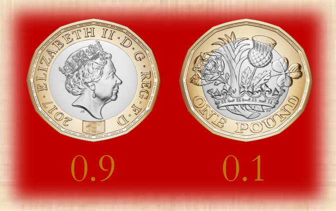

If she had instead learned the **red** model, she would have bet **heads** every period and earned an average payoff of 0.7 dollars!

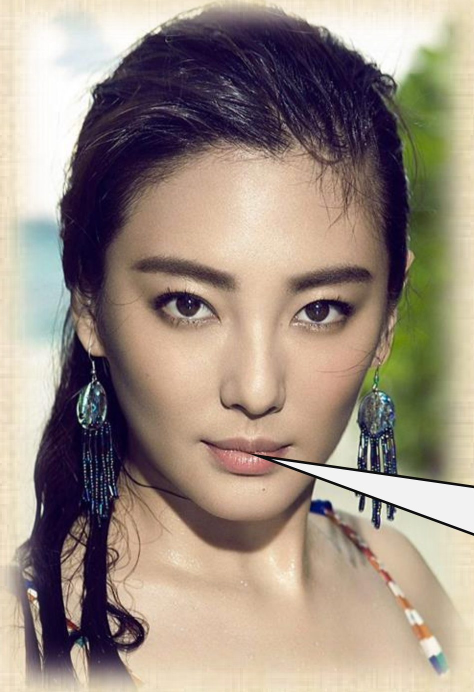

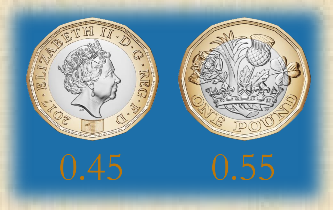

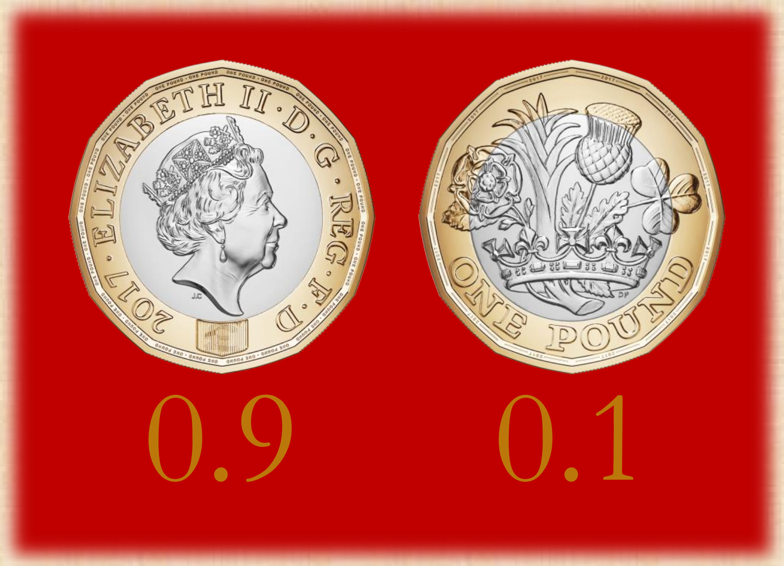

Koike is not a Bayesian.

Rather than update her prior probabilities over **red** and **blue** using standard likelihoods, she uses **generalized likelihoods** that depend on payoffs.

Let’s start with the **blue** model.

Recall that a believer in this model will bet on **tails**.

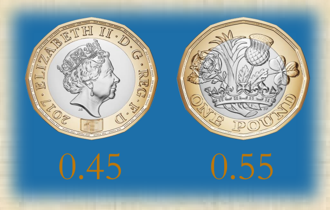

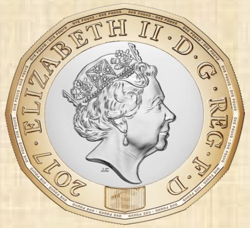

If she followed the **blue** model and **heads** arose, she would obtain a payoff of 0.

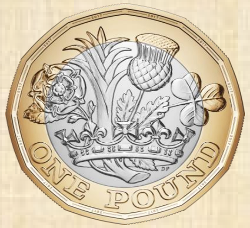

If she followed the **blue** model and **tails** arose, she would obtain a payoff of 1.

Take a sequence of observations, say

heads, tails, tails…

The likelihood of this sequence under the **blue** model is

$$(0.45)(0.55)(0.55)\ \ldots$$

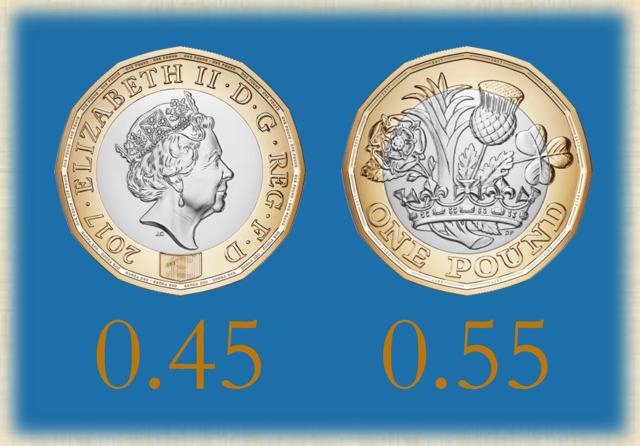

Take a sequence of observations, say

heads, tails, tails…

The **generalized likelihood** under the **blue** model is

$$\frac{e^{0}}{e^{0}+e^{1}} \cdot \frac{e^{1}}{e^{0}+e^{1}} \cdot \frac{e^{1}}{e^{0}+e^{1}} \cdots$$

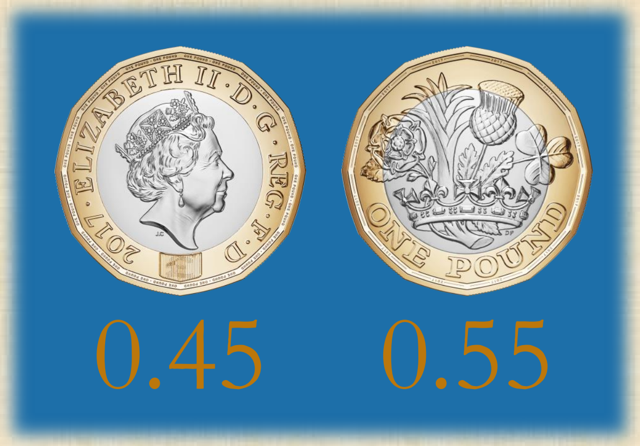

Exponents in numerator are the payoffs from the sequence **heads, tails, tails**… when the **blue** model is followed.

The denominator normalizes.

Do the same for the **red** model.

Now applying standard Bayesian learning using these generalized likelihoods…

…over time she comes to place all weight on the **red** model.

Hence she bets **heads** and earns an average payoff of 0.7 dollars.

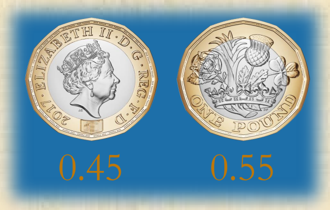

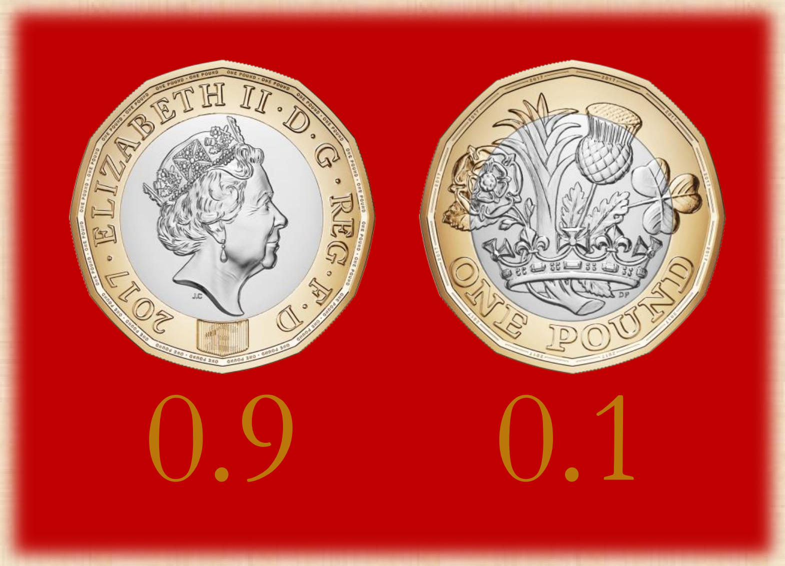

The model she learns is in some sense **wrong**, but in a pragmatic sense it works out just as well as if she had learnt the correct probability of **heads**.

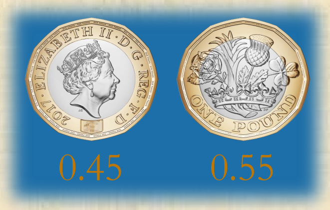

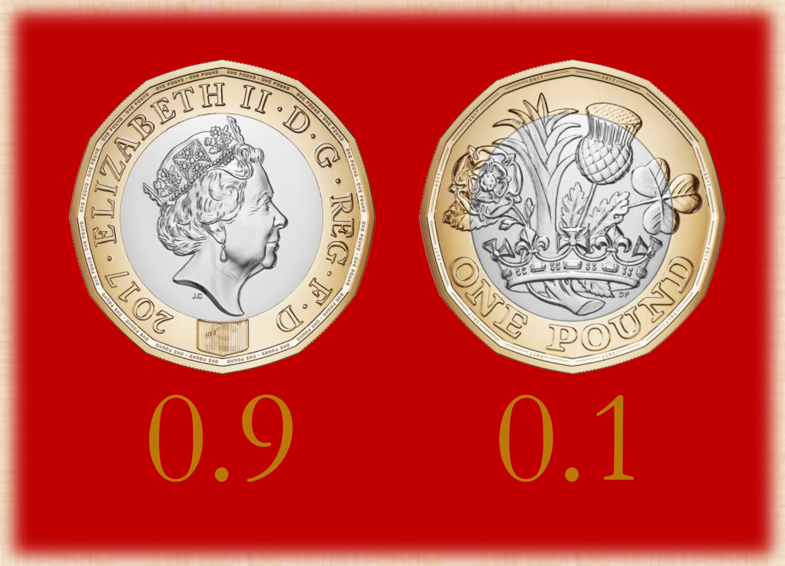

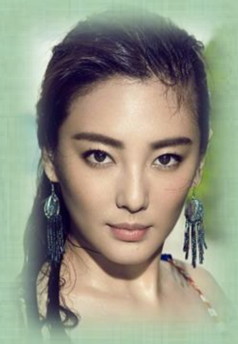

Bayesian learning

Axiomatically sound.

Learns about log-likelihoods, not payoffs.

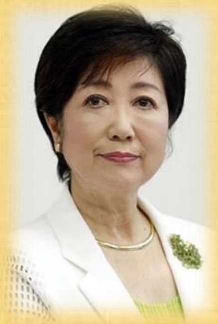

Generalized Bayesian learning

Pragmatically sound.

Learns about payoffs.

Note, if the model is correctly specified, so that $p(H) = 0.7$ is considered possible, then Bayesians learn the correct parameter.
Generalized Bayesians learn the correct parameter and all other parameters that lead them to bet the same way (i.e. **heads**).

Over time the Generalized Bayesians may outperform the standard Bayesians.

If payoffs are linked to fitness, Generalized Bayesians come to predominate in the population.

Over time the Generalized Bayesians may outperform the standard Bayesians.

If payoffs are linked to fitness, Generalized Bayesians come to predominate in the population.

Over time the Generalized Bayesians may outperform the standard Bayesians.

If payoffs are linked to fitness, Generalized Bayesians come to predominate in the population.

So far we have considered a one player decision problem.

What about games with more than one player?

What kind of behavior would these Generalized Bayesian, payoff optimizers play in equilibrium in a misspecified environment?

vs.

Answer is simple. Given the behavior of the other players, they should play according to whichever of their possible models of the world leads to the highest payoff.

Simple example.

Two players, **Alice** and **Betti**.

Betti

Alice

	b_1	b_2	b_3
a_1	$2, 2$	$1, 0$	$4, 1$
a_2	$1, 4$	$1, 5$	$3, 3$
a_3	$0, 1$	$0, 0$	$5, 5$

If **Alice** considers it impossible that **Betti** plays b_3 with any significant probability, then **Alice** will never play a_3.

If **Betti** considers it impossible that **Alice** plays a_2 with any significant probability, then **Betti** will never play b_2.

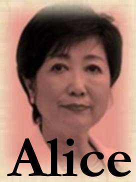

	b_1	b_2	b_3
a_1	$2, 2$	$1, 0$	$4, 1$
a_2	$1, 4$	$1, 5$	$3, 3$
a_3	$0, 1$	$0, 0$	$5, 5$

Consider the game in which we **exclude actions** that cannot be justified by any possibilities that the players consider.

Betti

Alice

	b_1	b_2	b_3
a_1	2, 2	1, 0	4, 1
a_2	1, 4	1, 5	3, 3
a_3	0, 1	0, 0	5, 5

Consider the game in which we **exclude actions** that cannot be justified by any possibilities that the players consider.

Find the **Nash equilibria**.

	b_1	~~b_2~~	b_3
a_1	**2, 2**	~~1, 0~~	4, 1
a_2	1, 4	~~1, 5~~	3, 3
~~a_3~~	~~0, 1~~	~~0, 0~~	~~5, 5~~

Consider the game in which we **exclude actions** that cannot be justified by any possibilities that the players consider.

Find the **Nash equilibria**.

This is a **misspecified Nash equilibrium (mNE)**.

Betti

Alice

	b_1	b_2	b_3
a_1	2, 2	1, 0	4, 1
a_2	1, 4	1, 5	3, 3
a_3	0, 1	0, 0	5, 5

At **mNE**, for any given player there are no beliefs in their model of the world that would justify them taking actions that would lead to their obtaining a higher payoff.

That is, there is **no profitable deviation**.

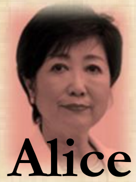

	b_1	b_2	b_3
a_1	2, 2	1, 0	4, 1
a_2	1, 4	1, 5	3, 3
a_3	0, 1	0, 0	5, 5

Note that a **NE** of the original game that is still feasible will remain a **NE** of the restricted game and will thus be a **mNE**.

This is the case with (a_1, b_1).

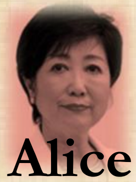

	b_1	~~b_2~~	b_3
a_1	**2, 2**	~~1, 0~~	4, 1
a_2	1, 4	~~1, 5~~	3, 3
~~a_3~~	~~0, 1~~	~~0, 0~~	~~5, 5~~

However, **mNE** need not be **NE** of the original game.

Consider the same game as before, but exchanging **Alice's** payoffs from a_1 and a_2.

Betti

Alice

	b_1	b_2	b_3
a_1	1, 2	1, 0	3, 1
a_2	2, 4	1, 5	4, 3
a_3	0, 1	0, 0	5, 5

Applying the same logic as before…

…(a_2, b_1) is an **NE** of the restricted game and thus a **mNE**.

However, (a_2, b_1) it is not a **NE** of the original game.

Betti

Alice

	b_1	b_2	b_3
a_1	$1, 2$	$1, 0$	$3, 1$
a_2	$2, 4$	$1, 5$	$4, 3$
a_3	$0, 1$	$0, 0$	$5, 5$

If generalized Bayesian learners converge to an equilibrium, then this equilibrium must be a mNE.

Betti

Alice

	b_1	b_2	b_3
a_1	1, 2	1, 0	3, 1
a_2	2, 4	1, 5	4, 3
a_3	0, 1	0, 0	5, 5

If a population converges to something other than mNE play, it will be vulnerable to invasion by generalized Bayesian learners.

Conversely, populations playing mNE are robust to invasion by types of learners who do not play mNE. *(all subject to T&Cs)*

Betti

Alice

	b_1	~~b_2~~	b_3
a_1	1, 2	~~1, 0~~	3, 1
a_2	**2, 4**	~~1, 5~~	4, 3
~~a_3~~	~~0, 1~~	~~0, 0~~	~~5, 5~~

More generally (and vaguely), the relationship between mNE and NE of the restricted game implies that arguments **in favour of/against** learning NE in correctly specified environments also work for mNE in misspecified environments.

Betti

Alice

	b_1	b_2	b_3
a_1	1, 2	1, 0	3, 1
a_2	2, 4	1, 5	4, 3
a_3	0, 1	0, 0	5, 5

Can regard beliefs about the world as justifications. Some are more **pragmatically useful** than others.

In this latter respect, **whether some beliefs are "truer" than other beliefs is incidental**.

This is something missed in correctly specified models, where maximum likelihood estimates and pragmatically useful beliefs coincide.

Thanks for listening!

Read the working paper at: **https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3473630**