From Positive to Zero to Negative Probability*

Adam Brandenburger

J.P. Valles Professor
Vice Dean for Graduate Education
NYU Stern School of Business

*Talk prepared for the Conference in Honor of the 60th Birthday of Sergei Artemov, renowned logician whom I am honored to call my friend

"Probability theory is nothing but common sense reduced to calculation."

-- Laplace: *Théorie Analytique des Probabilités*, 1812

Epistemic Game Theory

Conventional game theory says a model is a game matrix or tree

Epistemic Game Theory (EGT) says no: this is only a partial model

A full model consists of a matrix or tree together with a **belief structure** --- i.e., a space of possible beliefs, beliefs about beliefs, . . . , for each player

The beliefs are about what is uncertain --- both the structure of the game and the strategies chosen in the game

EGT respects the ‘trilogy’ of decision theory: strategies, payoffs, and probabilities

Zero Probability in Game Theory

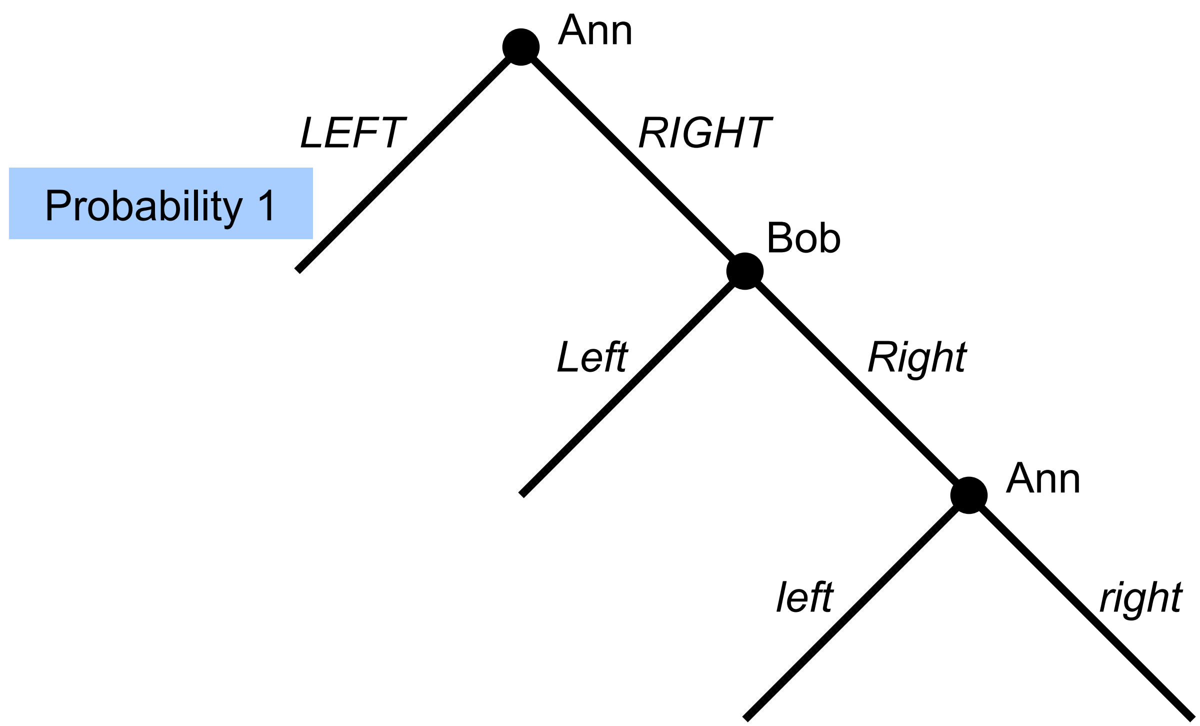

Suppose Bob assigns probability 1 to Ann's playing *LEFT*

Then, Bob's conditional probabilities on *RIGHT-left* vs. *RIGHT-right* are undefined

Of course, Bob's expected payoff is well-defined

Example cont'd

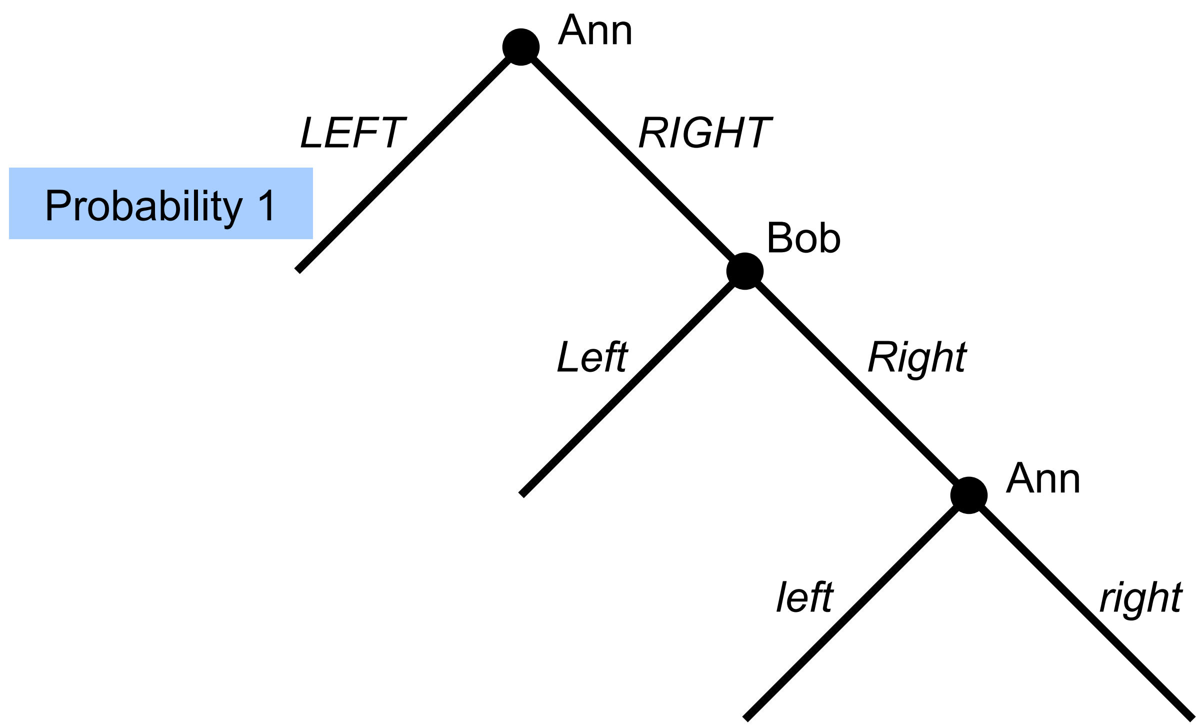

Now suppose Ann assigns positive probability to the event:

"Bob assigns probability 1 to Ann's playing *LEFT*"

Question:

What, in fact, should Ann play at her first node?!

This depends on how she thinks Bob would react to *RIGHT*

Chance vs. Control

What is **chance** from the point of view of one player (Bob) may be under the **control** of another player (Ann)

Probability-0 events cannot be neglected (even in a finite setting)

This is an intrinsically multi-player --- i.e., game-theoretic --- phenomenon

Conditioning on Probability-0 Events

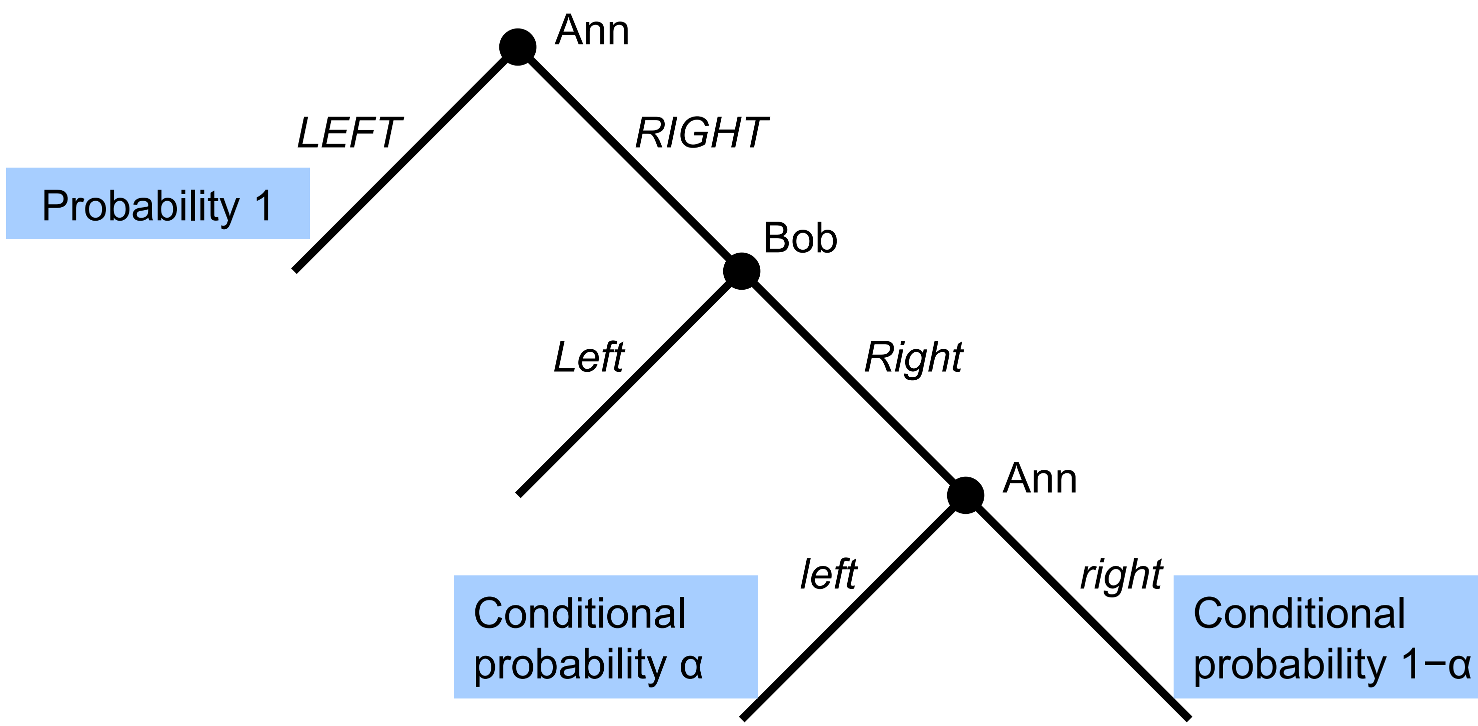

Ann needs a view as to how Bob would update his probabilities conditional on the probability-0 event that Ann chooses *RIGHT-left* or *RIGHT-right*

We need an extension to Kolmogorov probability theory

(Why not simply require all probabilities to be strictly positive?

We will see …)

Now, in the Matrix

Bob

		Left	*Right*
Ann	*Up*	2, 2	2, 2
	Down	1, 1	3, 3

Left is (weakly) **dominated** by *Right*

There is no full-support probability distribution on {*Up*, *Down*} which makes *Left* optimal

This is a general equivalence (Arrow-Barankin-Blackwell, 1953)

Question:

Should Ann put probability 0 or probability > 0 on *Left*?

An Extended Probability Theory

Ann possesses a **lexicographic probability system** (**LPS**), which is a sequence of probability measures (Blume, Brandenburger, and Dekel, 1991)

An LPS is used lexicographically in determining an optimal strategy:

1. pick those strategies that maximize expected payoff under the first probability measure
2. from this set, pick those strategies that maximize expected payoff under the second probability measure

etc.

Intuitively (and formally):

Bob's strategies that receive primary probability > 0
are infinitely more likely than
Bob's strategies that receive secondary probability > 0
are infinitely more likely than
…

We Solve Our Problem in the Tree, Too

Proposition (Blume, Brandenburger, and Dekel, 1991): *A strategy is admissible (undominated) if and only if there is a full-support LPS under which it is optimal.*

Proposition (Brandenburger, 2007): *A strategy is admissible if and only if for every tree that 'reduces' to the matrix, there is a conditional probability system (CPS) under which the strategy is optimal in that tree.*

A **CPS** (Rényi, 1955) specifies for each event in a given family of conditioning events, a probability measure that is proper (Blackwell and Ryll-Nardzewski, 1963) and obeys a chain rule

Rényi was interested in statistics and quantum mechanics (not game theory)

Battigalli and Siniscalchi (1999, 2002) developed and applied Rényi's theory to game trees

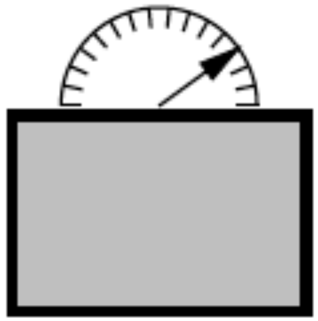

Alice chooses measurement setting a or a'

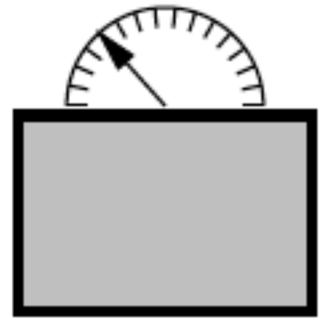

Bob chooses measurement setting b or b'

Here is an example of an **empirical model**:

	(0, 0)	**(1, 0)**	**(0, 1)**	**(1, 1)**
a b	1/2	0	0	1/2
a' b	3/8	1/8	1/8	3/8
a b'	3/8	1/8	1/8	3/8
a' b'	1/8	3/8	3/8	1/8

The Extended Model

	a	***a'***	***b***	***b'***
ω_0	0	0	0	0
ω_1	0	0	0	1
ω_2	0	0	1	0
⋮	⋮	⋮	⋮	⋮
ω_{15}	1	1	1	1

We look for a probability measure *p* on Ω that induces the empirical probabilities

In the classical domain, we can always find such a *p* (perform the joint measurements!)

In the quantum domain, we may not be able to (measurements may be incompatible!)

Introduce Negative Probabilities

Theorem (Abramsky and Brandenburger, 2011): *An empirical model is 'no signaling' if and only if there is an extended model with a signed probability measure that induces it.*

But, we overshoot quantum mechanics --- there are no-signaling empirical models with superquantum correlations (Popescu and Rohrlich, 1994)

	(0, 0)	**(1, 0)**	**(0, 1)**	**(1, 1)**
$a\ b$	1/2	0	0	1/2
$a^{I}\ b$	1/2	0	0	1/2
$a\ b^{I}$	1/2	0	0	1/2
$a^{I}\ b^{I}$	0	1/2	1/2	0

"Negative energies and probabilities should not be considered as nonsense. They are well-defined concepts mathematically, like a negative of money."

-- Dirac: "The Physical Interpretation of Quantum Mechanics," 1942