Title: Foundations and Interpretation of Quantum Theory - Lecture 7A

Date: Feb 04, 2010 02:30 PM

URL: http://pirsa.org/10020028

Abstract: After a review of the axiomatic formulation of quantum theory, the generalized operational structure of the theory will be introduced (including POVM measurements, sequential measurements, and CP maps). There will be an introduction to the orthodox (sometimes called Copenhagen) interpretation of quantum mechanics and the historical problems/issues/debates regarding that interpretation, in particular, the measurement problem and the EPR paradox, and a discussion of contemporary views on these topics. The majority of the course lectures will consist of guest lectures from international experts covering the various approaches to the interpretation of quantum theory (in particular, many-worlds, de Broglie-Bohm, consistent/decoherent histories, and statistical/epistemic interpretations, as time permits) and fundamental properties and tests of quantum theory (such as entanglement and experimental tests of Bell inequalities, contextuality, macroscopic quantum phenomena, and the problem of quantum gravity, as time permits).

Probability Course, Day Two

Last time, I spent a lot of time trying to summarize, explain, and differentiate the frequentist and Bayesian viewpoints on probability. (The fact that some of you are itching to raise your hand and ask "Don't you mean physical versus evidentiary?" merely warms my heart; if that irritates you, I've done my job.) Today, I want to turn around and start bridging that gap — bringing the two viewpoints back into some semblance of consistency.

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

* Reading for next class

http://plato.stanford.edu/entries/bell-theorem/

by Shimony of (CHSH)

* Assignment #1 : (a) Sent out by email on UN-ACE

(b) due Feb 11th

Probability Course, Day Two

Last time, I spent a lot of time trying to summarize, explain, and differentiate the frequentist and Bayesian viewpoints on probability. (The fact that some of you are itching to raise your hand and ask "Don't you mean physical versus evidentiary?" merely warms my heart; if that irritates you, I've done my job.) Today, I want to turn around and start bridging that gap — bringing the two viewpoints back into some semblance of consistency.

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:41 PM Robin Blume-Kohout

Robin's Notes - An ongoing quest for coherence

file:///Users/rbk/Documents/RBKwiki/index.html

Probability Course, Day Two

Last time, I spent a lot of time trying to summarize, explain, and differentiate the frequentist and Bayesian viewpoints on probability. (The fact that some of you are itching to raise your hand and ask "Don't you mean physical versus evidentiary?" merely warms my heart; if that irritates you, I've done my job.) Today, I want to turn around and start bridging that gap — bringing the two viewpoints back into some semblance of consistency.

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:41 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

file:///Users/rbk/Documents/RBKwiki/index.html

Probability Course, Day Two

Last time, I spent a lot of time trying to summarize, explain, and differentiate the frequentist and Bayesian viewpoints on probability. (The fact that some of you are itching to raise your hand and ask "Don't you mean physical versus evidentiary?" merely warms my heart; if that irritates you, I've done my job.) Today, I want to turn around and start bridging that gap — bringing the two viewpoints back into some semblance of consistency.

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

file:///Users/rbk/Documents/RBKwiki/index.html

iGoogle TinyURL! Proxyify Add to Wish List Kraft's inequality – ... UWaterloo E-Journals Notebook Picasa Web Albums ... Readability EXIF

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

Probability Course, Day Two

+ - Ø ¤ ×

Last time, I spent a lot of time trying to summarize, explain, and differentiate the frequentist and Bayesian viewpoints on probability. (The fact that some of you are itching to raise your hand and ask "Don't you mean physical versus evidentiary?" merely warms my heart; if that irritates you, I've done my job.) Today, I want to turn around and start bridging that gap — bringing the two viewpoints back into some semblance of consistency.

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

file:///Users/rbk/Documents/RBKwiki/index.html

any pages is princip

iGoogle TinyURL! Proxyify Add to Wish List Kraft's inequality – ... UWaterloo E-Journals Notebook Picasa Web Albums ... Readability EXIF

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

I want to turn around and start bridging that gap — bringing the two viewpoints back into some semblance of consistency.

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

file:///Users/rbk/Documents/RBKwiki/index.html

any pages is princip

iGoogle TinyURL! Proxyify Add to Wish List Kraft's inequality – ... UWaterloo E-Journals Notebook Picasa Web Albums ... Readability EXIF

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

Toolbars
✓ Status Bar
Sidebar
Stop ⌘.
Reload ⌘R
Zoom
Page Style
Character Encoding
Page Source ⌘U

…tes – An ongoing quest for coherence

…/RBKwiki/index.html

iGoogle TinyURL! Proxyify … UWaterloo E-Journals Notebook Picasa Web Albums … Readability EXIF

Robin's Notes – An o… | …obability axioms – … | Quantum Theory – | Principia Mathematic…

You're probabl … point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

Toolbars
✓ Status Bar
Sidebar
Stop ⌘.
Reload ⌘R
Zoom
Page Style
Character Encoding
Page Source ⌘U

✓ Navigation Toolbar
✓ Bookmarks Toolbar
Customize...

coherence

any pages is princip

iGoogle TinyURL! Proxyify otebook Picasa Web Albums ... Readability EXIF

Robin's Notes – An o... robability axioms – ... Quantum Theory – Principia Mathematic...

You're probabl point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

tes – An ongoing quest for coherence

s/RBKwiki/index.html

any pages is princip

iGoogle TinyURL! Proxyify ... UWaterloo E-Journals Notebook Picasa Web Albums ... Readability EXIF

Robin's Notes – An o... | robability axioms – ... | Quantum Theory – | Principia Mathematic...

Toolbars
✓ Status Bar
Sidebar

Stop ⌘.
Reload ⌘R

Zoom
Page Style
Character Encoding

Page Source ⌘U

Zoom In ⌘+
Zoom Out ⌘–

Reset ⌘0

Zoom Text Only

You're probabl[...] earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

Toolbars
✓ Status Bar
Sidebar
Stop ⌘.
Reload ⌘R
Zoom
Page Style
Character Encoding
Page Source ⌘U

tes – An ongoing quest for coherence

/RBKwiki/index.html

any pages is princip

iGoogle TinyURL! Proxyify ... UWaterloo E-Journals Notebook Picasa Web Albums ... Readability EXIF

Robin's Notes – An o... | robability axioms – ... | Quantum Theory – | Principia Mathematic...

You're probabl[...] point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:43 PM Robin Blume-Kohout

tes – An ongoing quest for coherence

/RBKwiki/index.html

any pages is princip

iGoogle TinyURL! Proxyify ... UWaterloo E-Journals Notebook Picasa Web Albums ... Readability EXIF

Robin's Notes – An o... robability axioms – ... Quantum Theory – Principia Mathematic...

Toolbars
✓ Status Bar
Sidebar

Stop ⌘.
Reload ⌘R

Zoom
Page Style
Character Encoding

Page Source ⌘U

No Style
✓ Basic Page Style

You're probabl point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper

You're probably wondering, at this point, why on earth I want to reconcile frequentist and Bayesian probability, given that I spent an hour stomping vigorously on the corpse of frequentism last time. By way of answer, here is an analogy: suppose that this class was called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper truths) about the world we live in. The Principle of Least Action is really important, even if Feynman sort of showed us that it's "just" a semiclassical manifestation of path integration in quantum theory. In almost every regime where it applies, classical mechanics is enormously simpler to calculate than quantum mechanics. Moreover, if you believe in hidden variables, then you believe that there is a hidden classical mechanical explanation *underneath* quantum mechanics, in which case Feynman's path integral formulation is itself just a complicated manifestation of the Principle of Least Action!

In other words, it just might be turtles all the way down.

Now, (a) I'm not here to talk about quantum mechanics, and (b) don't take that as a particularly strong analogy. Frequentism is not inherently associated with hidden variables or classical mechanics, nor is Bayesianism inherently associated with quantum mechanics. There's a correlations, perhaps, but it's sociological — people with *realist* inclinations tend to like (i) physical probability, and (ii) hidden variables. Anti-realists like myself tend to view both as naive, and we believe that we're forced into other, counterintuitive, viewpoints. Operationalists (and most working quantum mechanics are operationalists at least some of the time) go either way, and usually believe that their viewpoint (whichever it is) is the natural one for all operationalists. E.g., Joseph believes that operationalism is a natural fit

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:48 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper truths) about the world we live in. The Principle of Least Action is really important, even if Feynman sort of showed us that it's "just" a semiclassical manifestation of path integration in quantum theory. In almost every regime where it applies, classical mechanics is enormously simpler to calculate than quantum mechanics. Moreover, if you believe in hidden variables, then you believe that there is a hidden classical mechanical explanation *underneath* quantum mechanics, in which case Feynman's path integral formulation is itself just a complicated manifestation of the Principle of Least called "Mechanics," and I had spent the first period explaining in detail why classical mechanics is wrong, and quantum mechanics (though highly counterintuitive) is the only known alternative that *isn't* wrong (even though we don't actually understand it, and find it deeply unsatisfying at an ontological, interpretational level). Given what you know, would you be particularly surprised if I came back the next day and started teaching you Newton's Laws? What if I rephrased them as "Newton's Really Reliable Rules"?

The point, which I hope is obvious, is that classical mechanics is a hell of a good approximation in all sorts of circumstances — and, moreover, it reflects some higher truths (as opposed to deeper truths) about the world we live in. The Principle of Least Action is really important, even if Feynman sort of showed us that it's "just" a semiclassical manifestation of path integration in quantum theory. In almost every regime where it applies, classical mechanics is enormously simpler to calculate than quantum mechanics. Moreover, if you believe in hidden variables, then you believe that there is a hidden classical mechanical explanation *underneath* quantum mechanics, in which case Feynman's path integral

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:49 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

formulation is itself just a complicated manifestation of the Principle of Least Action!

In other words, it just might be turtles all the way down.

Now, (a) I'm not here to talk about quantum mechanics, and (b) don't take that as a particularly strong analogy. Frequentism is not inherently associated with hidden variables or classical mechanics, nor is Bayesianism inherently associated with quantum mechanics. There's a correlations, perhaps, but it's sociological — people with *realist* inclinations tend to like (i) physical probability, and (ii) hidden variables. Anti-realists like myself tend to view both as naive, and we believe that we're forced into other, counterintuitive, viewpoints. Operationalists (and most working quantum mechanics are operationalists at least some of the time) go either way, and usually believe that their viewpoint (whichever it is) is the natural one for all operationalists. E.g., Joseph believes that operationalism is a natural fit for frequentism, and I believe it's a natural fit for Bayesianism. We're both wrong, of course.

So, returning to my original point: frequentist views on probability are at least as useful as classical mechanics. I believe that, like classical mechanics, they are *not* suitable as a foundation — they fail to actually give a generally valid definition of one for all operationalists. E.g., Joseph believes that operationalism is a natural fit for frequentism, and I believe it's a natural fit for Bayesianism. We're both wrong, of course.

So, returning to my original point: frequentist views on probability are at least as useful as classical mechanics. I believe that, like classical mechanics, they are *not* suitable as a foundation — they fail to actually give a generally valid definition of what probability *is*, which was the entire point. Remember, please, that in mathematics and physics, the "most true" definition of something is frequently not a terribly intuitive one — it's just supposed to be the simplest definition that **works absolutely**.

So, for instance, if you've studied formal mathematics, then you'll recall that the natural numbers are defined in a rather awkward set-theoretic way,

$$0 = \{\emptyset\},\ \ 1 = \{0, \emptyset\},\ \ 2 = \{0, 1, \{\emptyset\}\}\ldots$$

and any sane 3rd grader might ask, "Oh, for crying out loud! Numbers are obvious — can't we just define N as a pile of N bottle caps or something?" But, of course, Bertrand Russell replies, "But then the natural numbers don't exist unless bottle caps do!" and you end up with 500 pages of *Principia Mathematica* justifying this counterintuitively complicated-looking *definition* of the natural numbers. Because anything simpler turns out to have loopholes.

Of course, anybody can *use* the natural numbers without reading *Principia Mathematica*, and (thank god!) we can use the real numbers without referring to Dedekind cuts. But, to quote a recent article by Lucien Hardy and Rob Spekkens,

> But, just as understanding how to drive an automobile is different from
> understanding how it works or how to fix it should it break down, so too is there a difference between understanding how to use quantum theory and understanding what it means.

The same is true of probability theory. On the other hand, it's a poor mechanic indeed who doesn't know how to drive! So let's take probability theory for a spin.

The Axiomatic Approach

Kolmogorov decided that he didn't care about what probability actually was. He

Firefox File Edit View History Bookmarks Tools Window Help Thu 2:52 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

counterintuitively complicated looking *definition* of the natural numbers. Because anything simpler turns out to have loopholes.

Of course, anybody can *use* the natural numbers without reading *Principia Mathematica*, and (thank god!) we can use the real numbers without referring to Dedekind cuts. But, to quote a recent article by Lucien Hardy and Rob Spekkens,

> But, just as understanding how to drive an automobile is different from
> understanding how it works or how to fix it should it break
> down, so too is there a difference between understanding
> how to use quantum theory and understanding what it
> means.

The same is true of probability theory. On the other hand, it's a poor mechanic indeed who doesn't know how to drive! So let's take probability theory for a spin.

The Axiomatic Approach

The Axiomatic Approach

Kolmogorov decided that he didn't care about what probability actually was. He was a mathematician, and mathematicians could care less about the physical world. They just need some axioms that tell them how stuff *works*, and they're good to go. So Kolmogorov established axioms for probability theory.

First, he had to establish a framework (within set theory) in which probability could be defined.

1. Let there be a set Ω whose elements are called *elementary events*. This is our *sample space*.
2. Let Σ be a set containing subsets of Ω. These are *events* — i.e., things that could be observed.
3. Let $Pr : \Sigma \to [0 \ldots 1]$ be a function that assigns a real number to each event. This is our *probability measure*.

Now, the triple (Ω, Σ, Pr) is a *probability space* if and only if it satisfies the following three "Kolmogorov Axioms".

The Axiomatic Approach

Kolmogorov decided that he didn't care about what probability actually was. He was a mathematician, and mathematicians could care less about the physical world. They just need some axioms that tell them how stuff *works*, and they're good to go. So Kolmogorov established axioms for probability theory.

First, he had to establish a framework (within set theory) in which probability could be defined.

1. Let there be a set Ω whose elements are called *elementary events*. This is our *sample space*.
2. Let Σ be a set containing subsets of Ω. These are *events* — i.e., things that could be observed.
3. Let $Pr : \Sigma \to [0 \ldots 1]$ be a function that assigns a real number to each event. This is our *probability measure*.

Now, the triple (Ω, Σ, Pr) is a *probability space* if and only if it satisfies the following three "Kolmogorov Axioms".

Now, the triple (Ω, Σ, Pr) is a *probability space* if and only if it satisfies the following three "Kolmogorov Axioms".

1. For all $E \in \Sigma$, $Pr(E) \geq 0$
2. $Pr(\Omega) = 1$;
3. $Pr\left(\bigcup_k \{E_k\}\right) = \sum_k Pr(E_k)$ where $E_k \in \Sigma$ and $E_i \cup E_j = \emptyset \; \forall \, i, j$.

In plain language, these axioms are:

1. Every probability is non-negative.
2. *Something* has to happen — so the set of all elementary events has probability 1.
3. If the events $\{E_k\}$ are disjoint, then the probability of their union is the sum of their probabilities.

The 3rd axiom has the most consequences; two special cases are:

- The empty set has $Pr(\emptyset) = 0$.
- If we have events $E, F \in \Sigma$ such that $E \subseteq F$, then $p(E) \leq p(F)$.

All the rules of probability theory — in particular, its **coherence** — can be derived from these axioms. However, probability "is" nothing but a *measure* — i.e., a function that assigns real numbers to subsets of Ω. No connection to physics or real life is made. However, we can do calculations all day within this framework, interpreting the results however we please.

Note that Kolmogorov's axioms (a) are a bit more complex than you might think is necessary (e.g., "Can't we just assign probabilities to elementary events, and dispense with Σ?"), and (b) allow some surprising and counterintuitive possibilities. For instance:

1. Σ may not contain any events corresponding to single elementary events! This is why we can't just assign probabilities to elementary events.
2. If it does, it may be that every elementary event has $Pr(x) = 0$... if there are uncountably many elementary events.
3. $Pr(E) = 0$ does not imply that E cannot happen.
4. Some sample spaces have subsets that *cannot* appear in Σ — i.e., are "non-measurable".

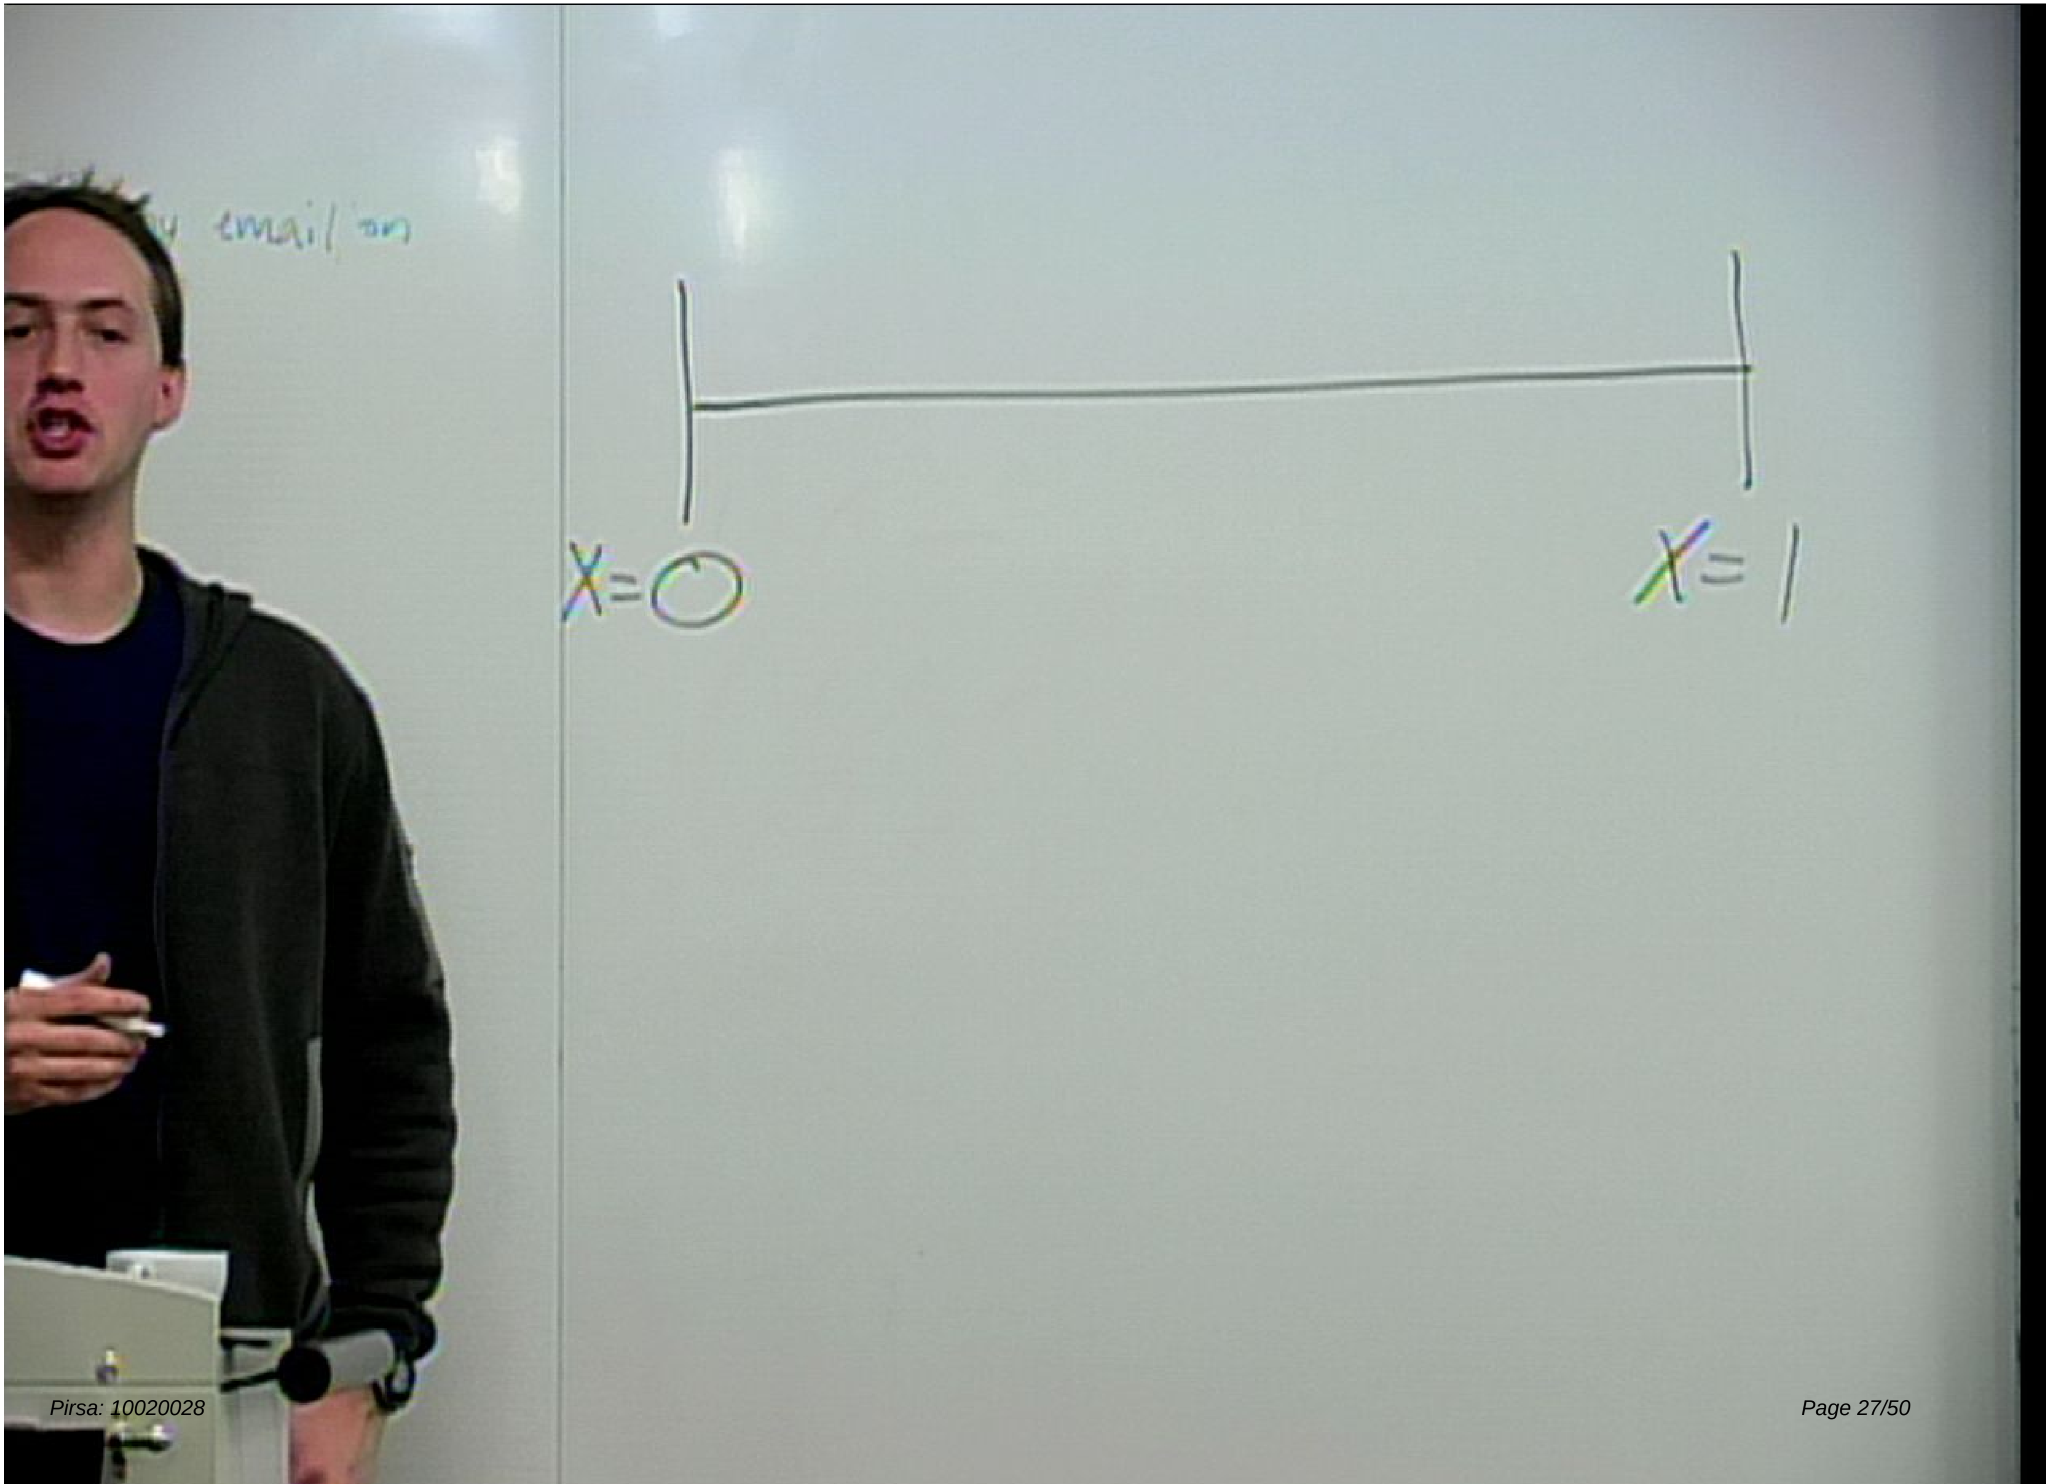
X=0
X=1

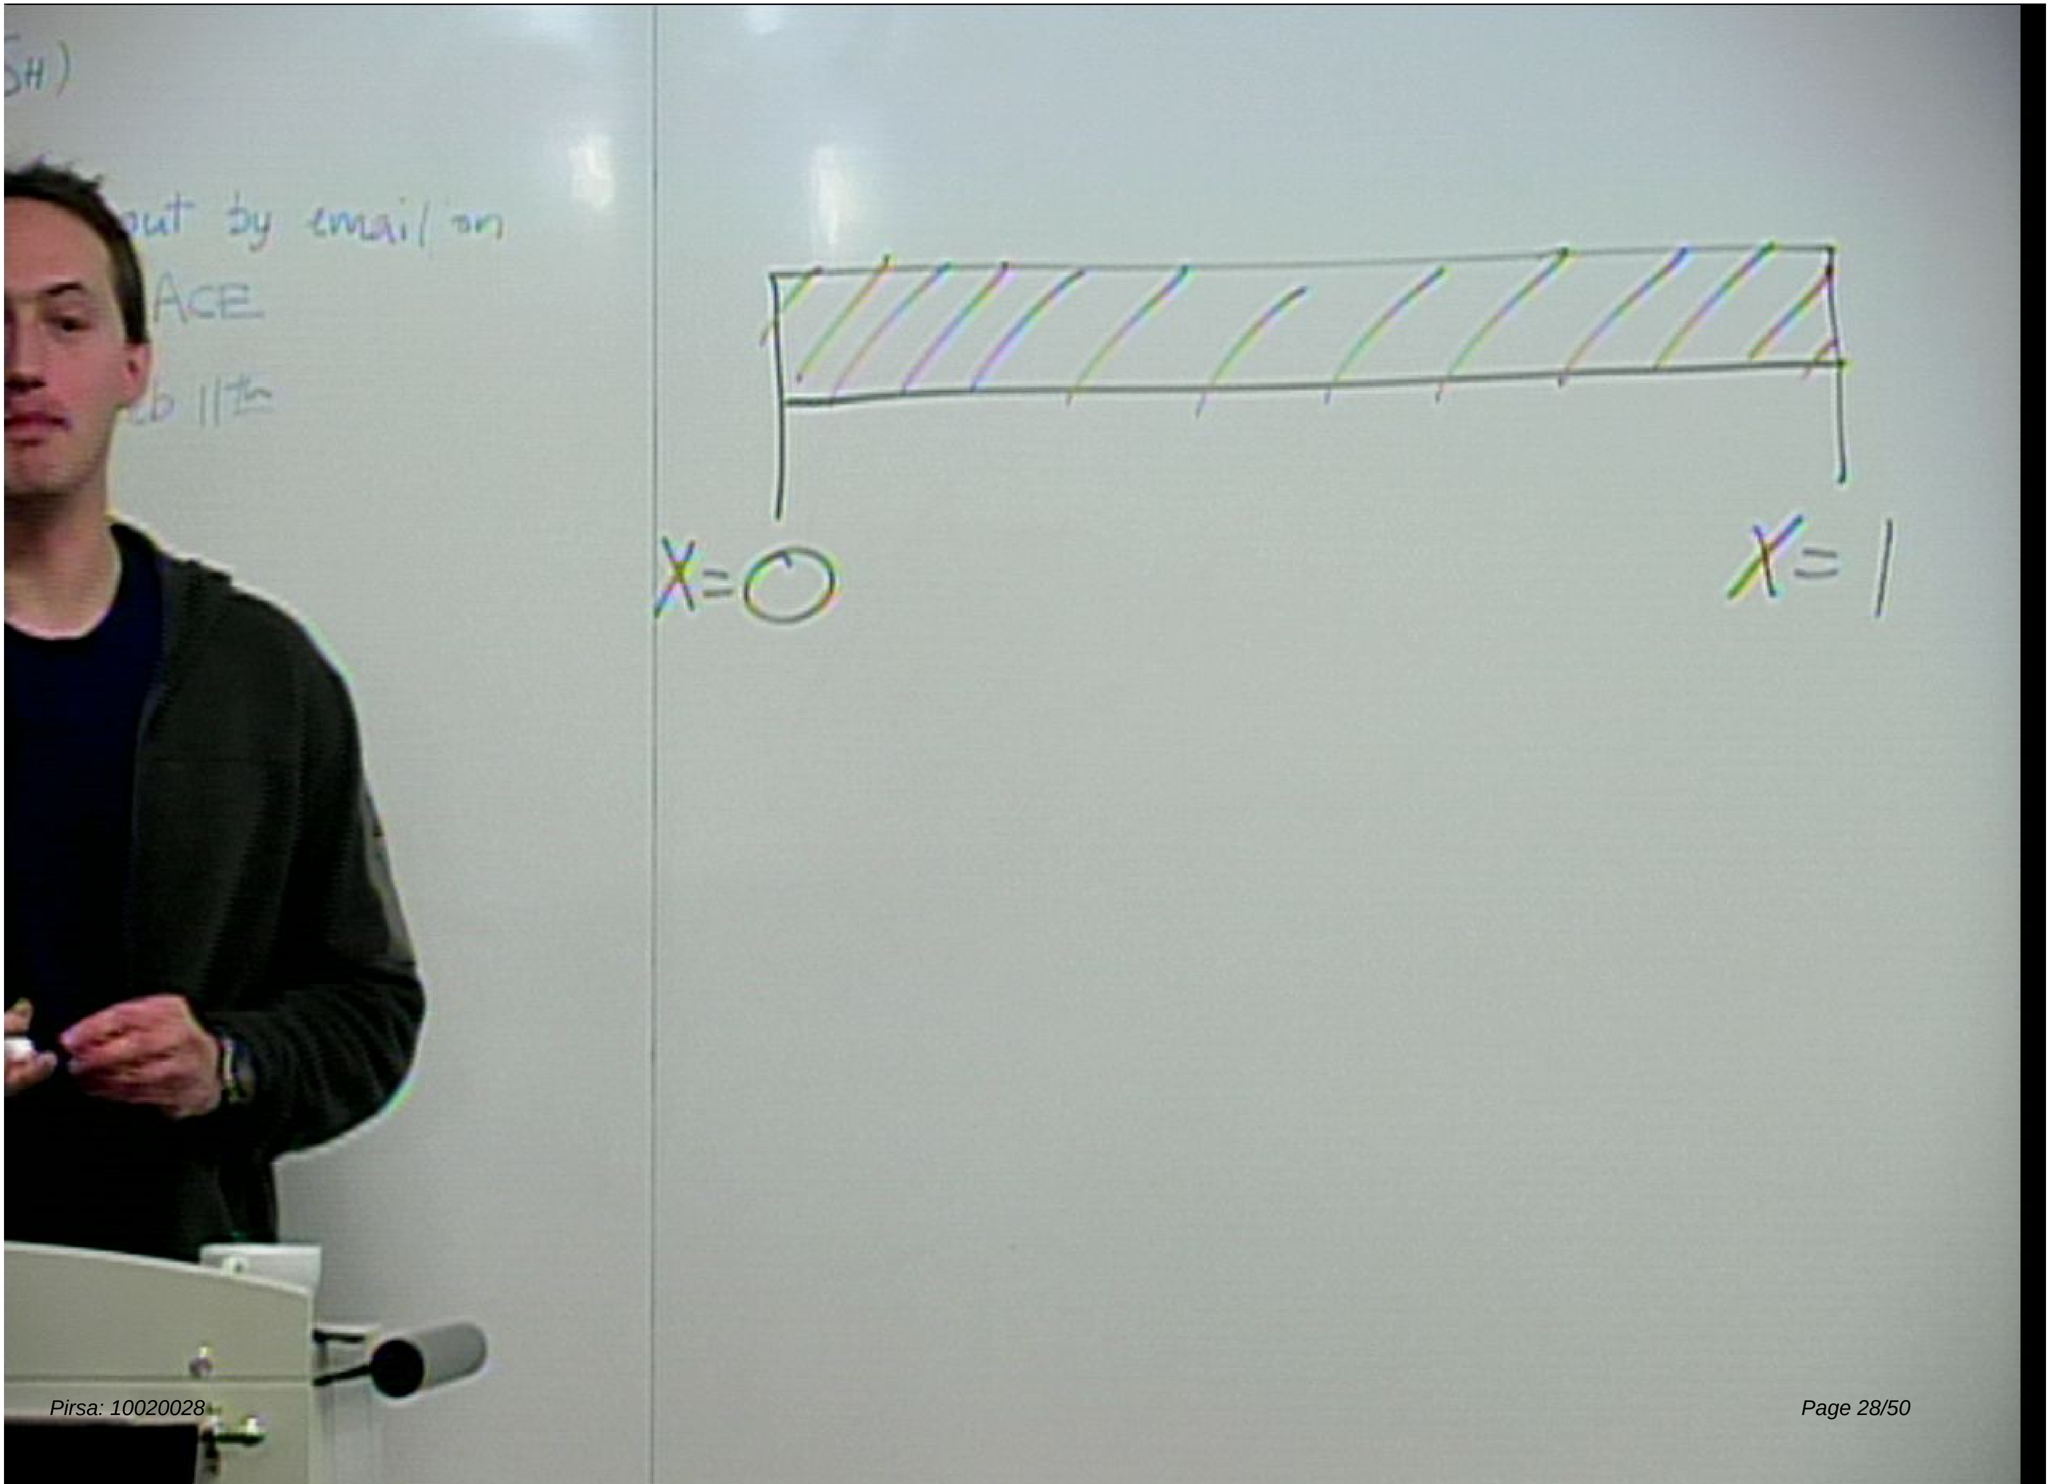
out by email on
ACE
X = 0
X = 1

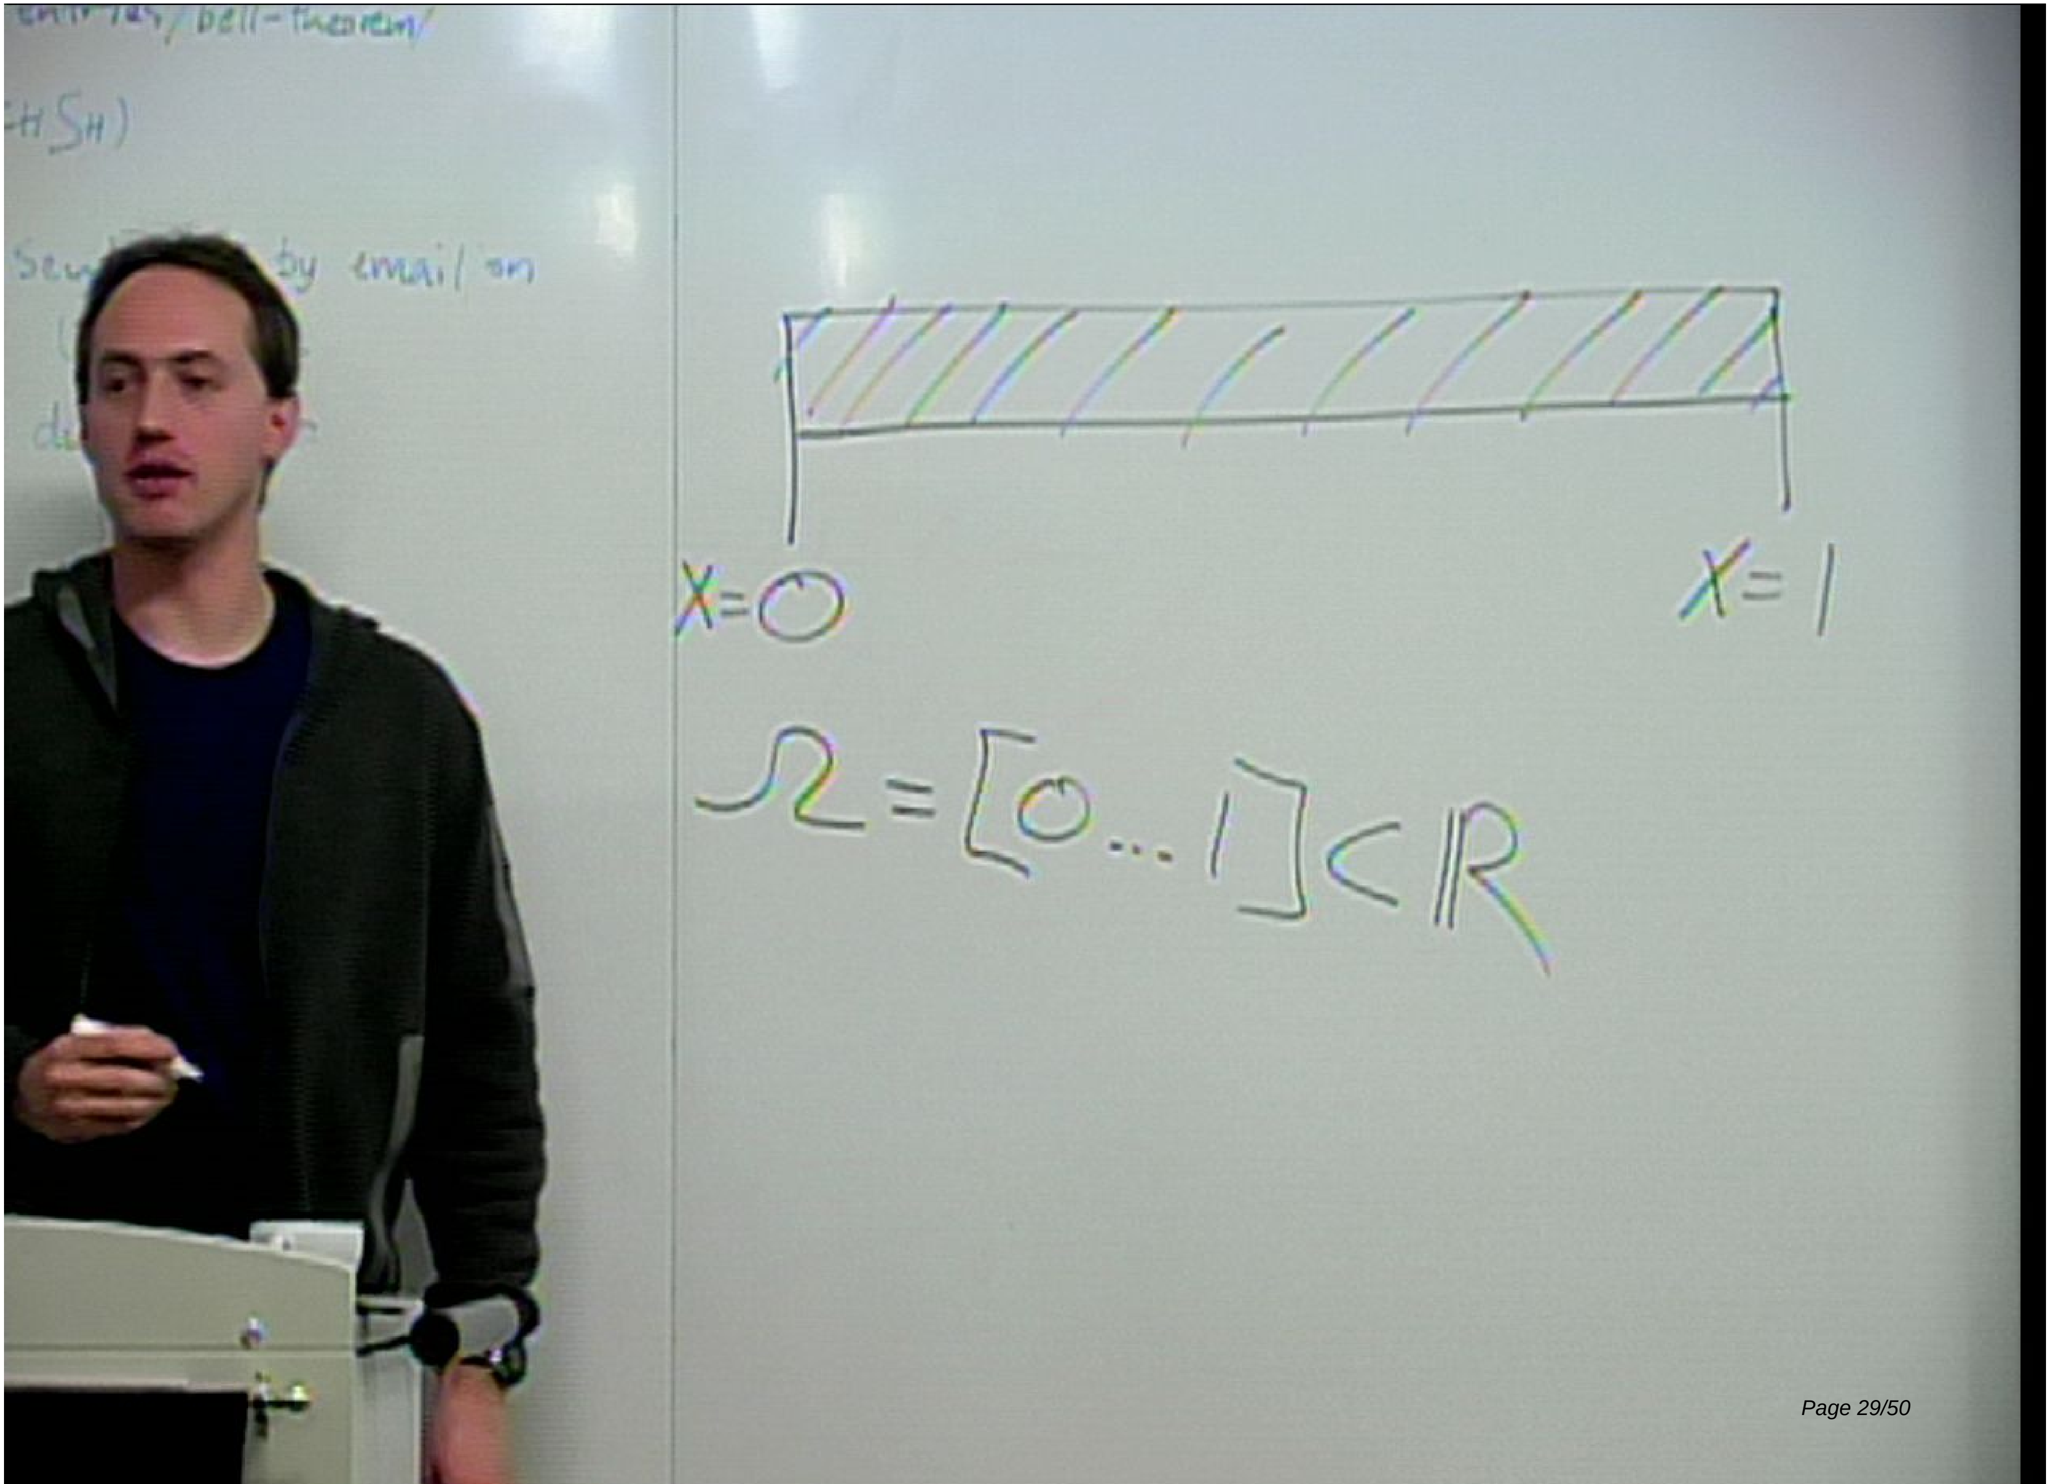
bell-theorem/
by email
X=0
X=1
Ω=[0...1]⊂R

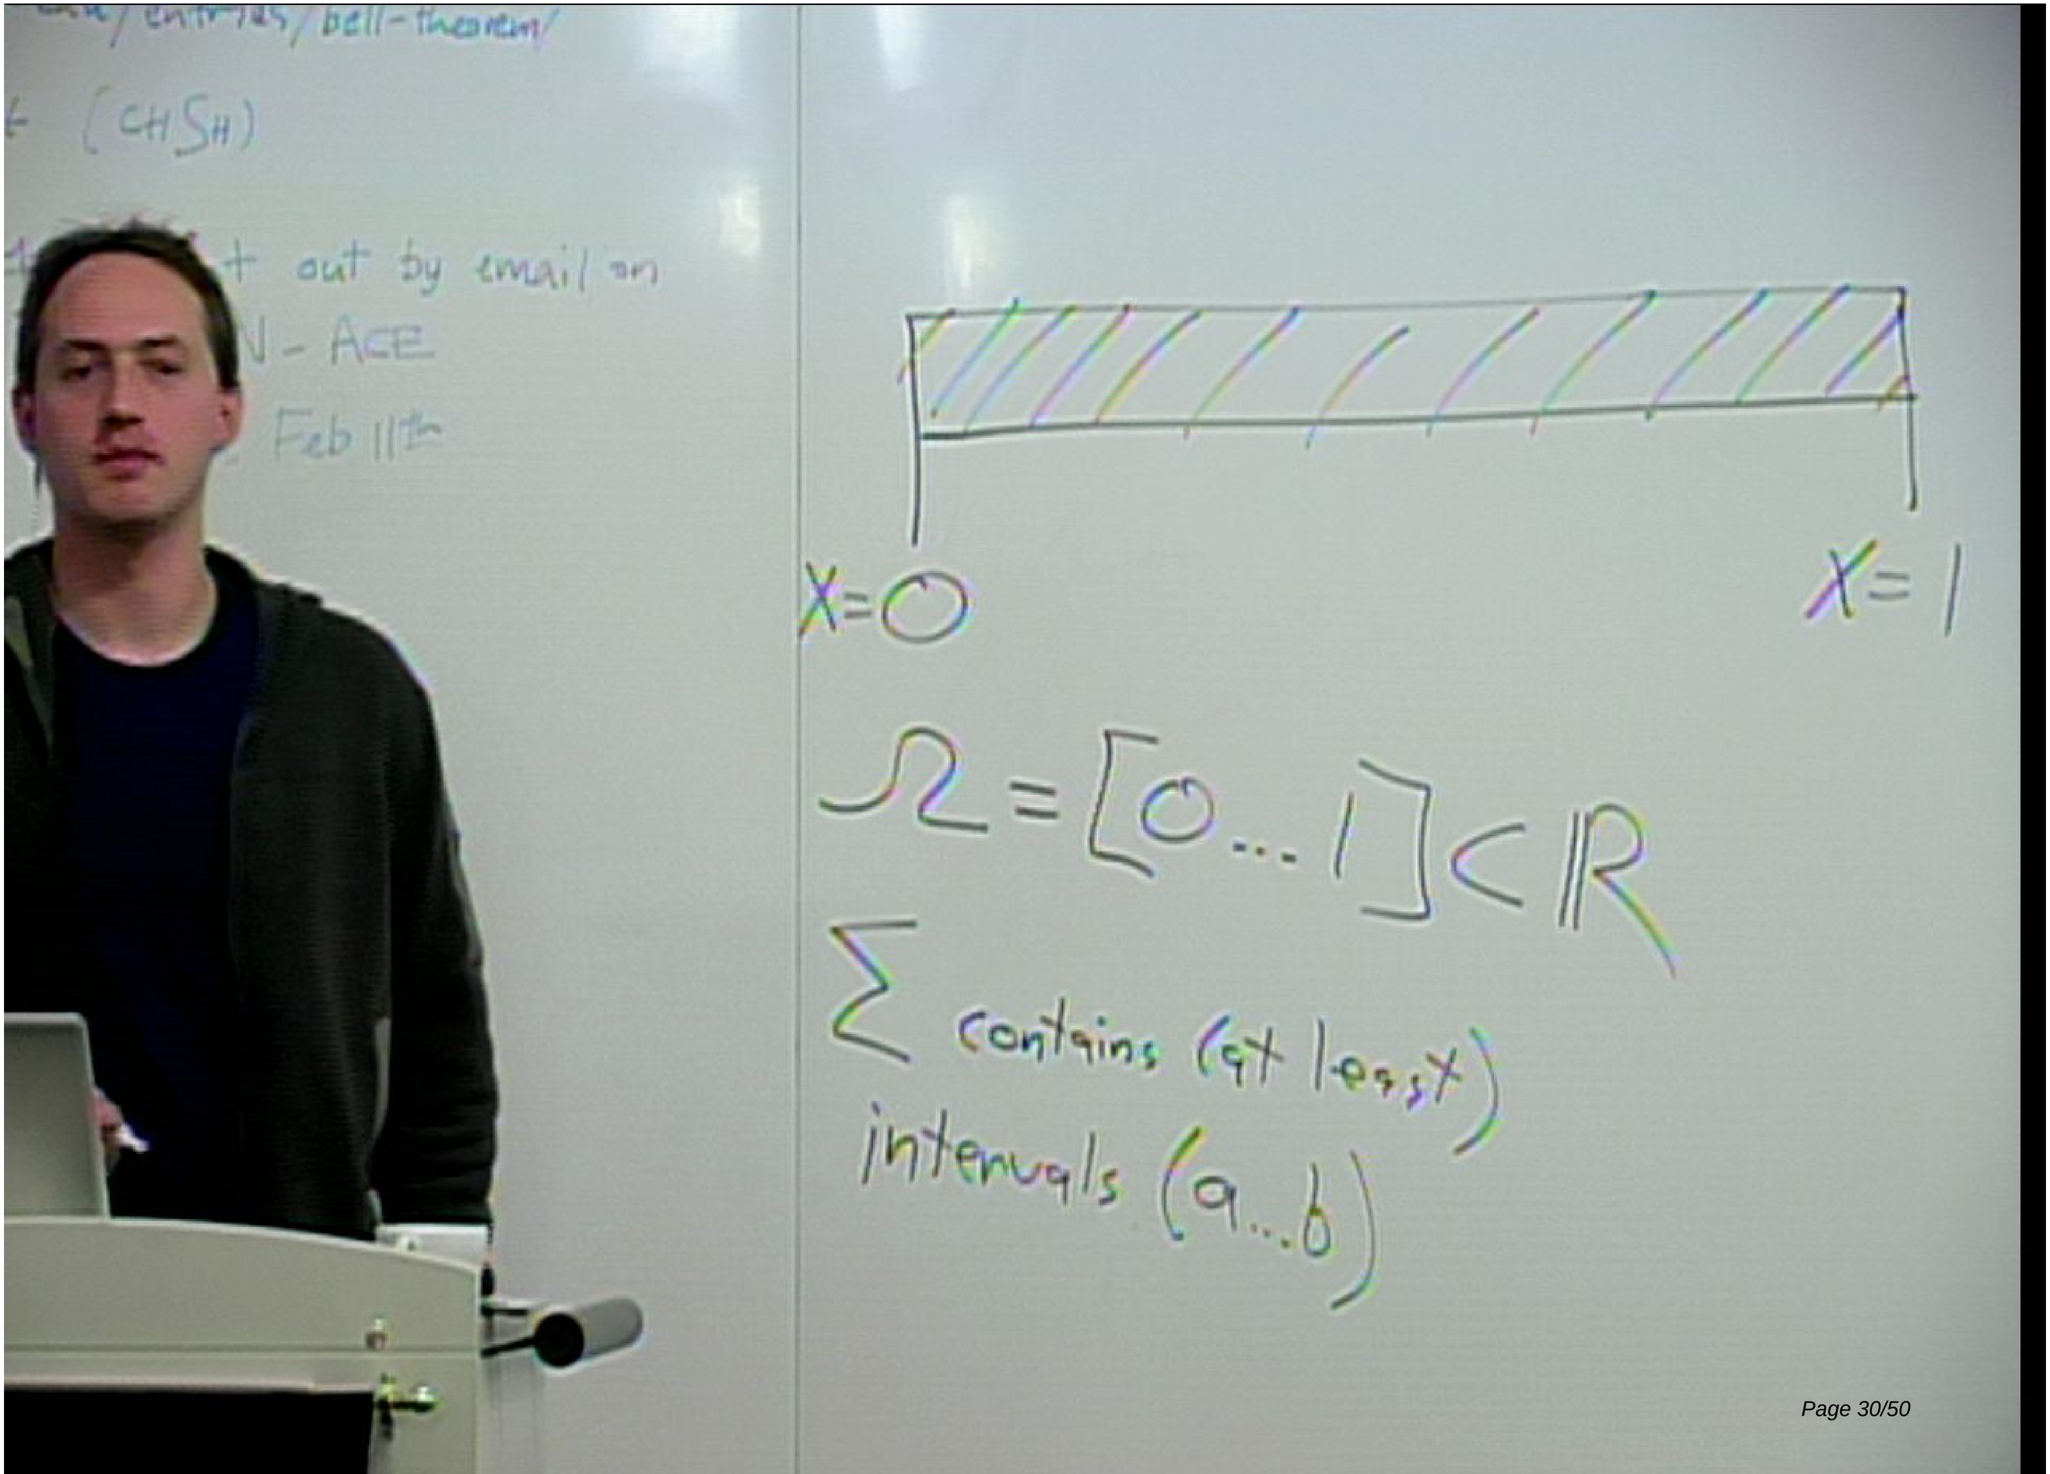
/entries/bell-theorem/
(CHSH)
out by email on
N - ACE
Feb 11th
X=0
X=1
Ω = [0...1] ⊂ ℝ
Σ contains (at least)
intervals (a...b)

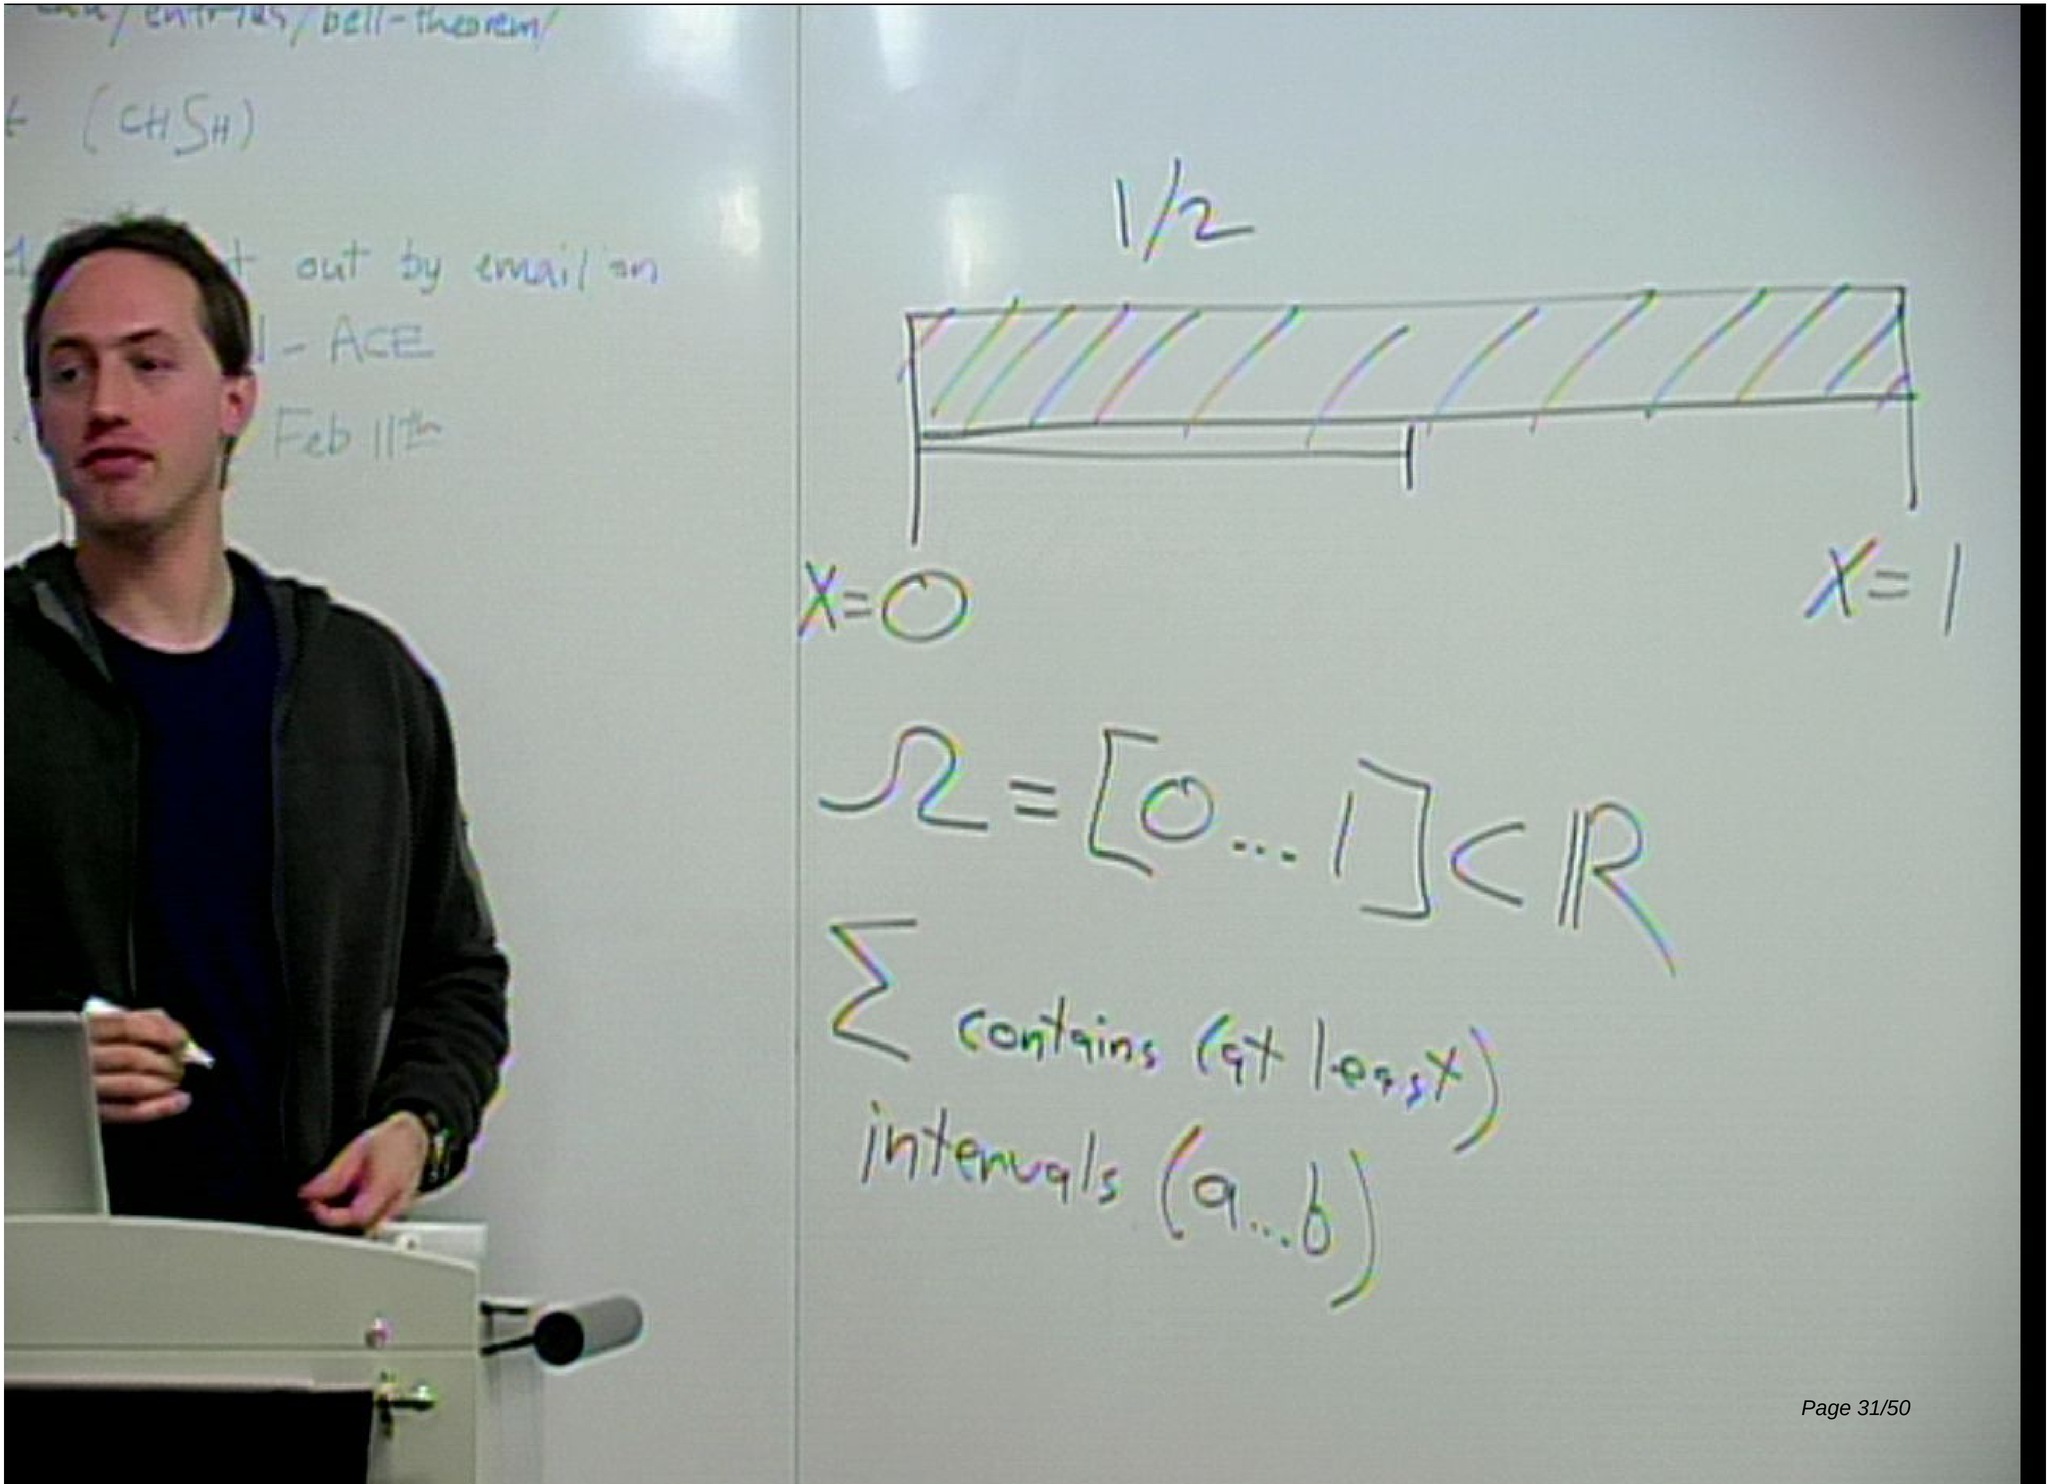
/entries/bell-theorem/
(CHSH)
out by email on
- ACE
Feb 11th
1/2
X=0
X=1
Ω=[0...1]⊂R
Σ contains (at least x)
intervals (a...b)

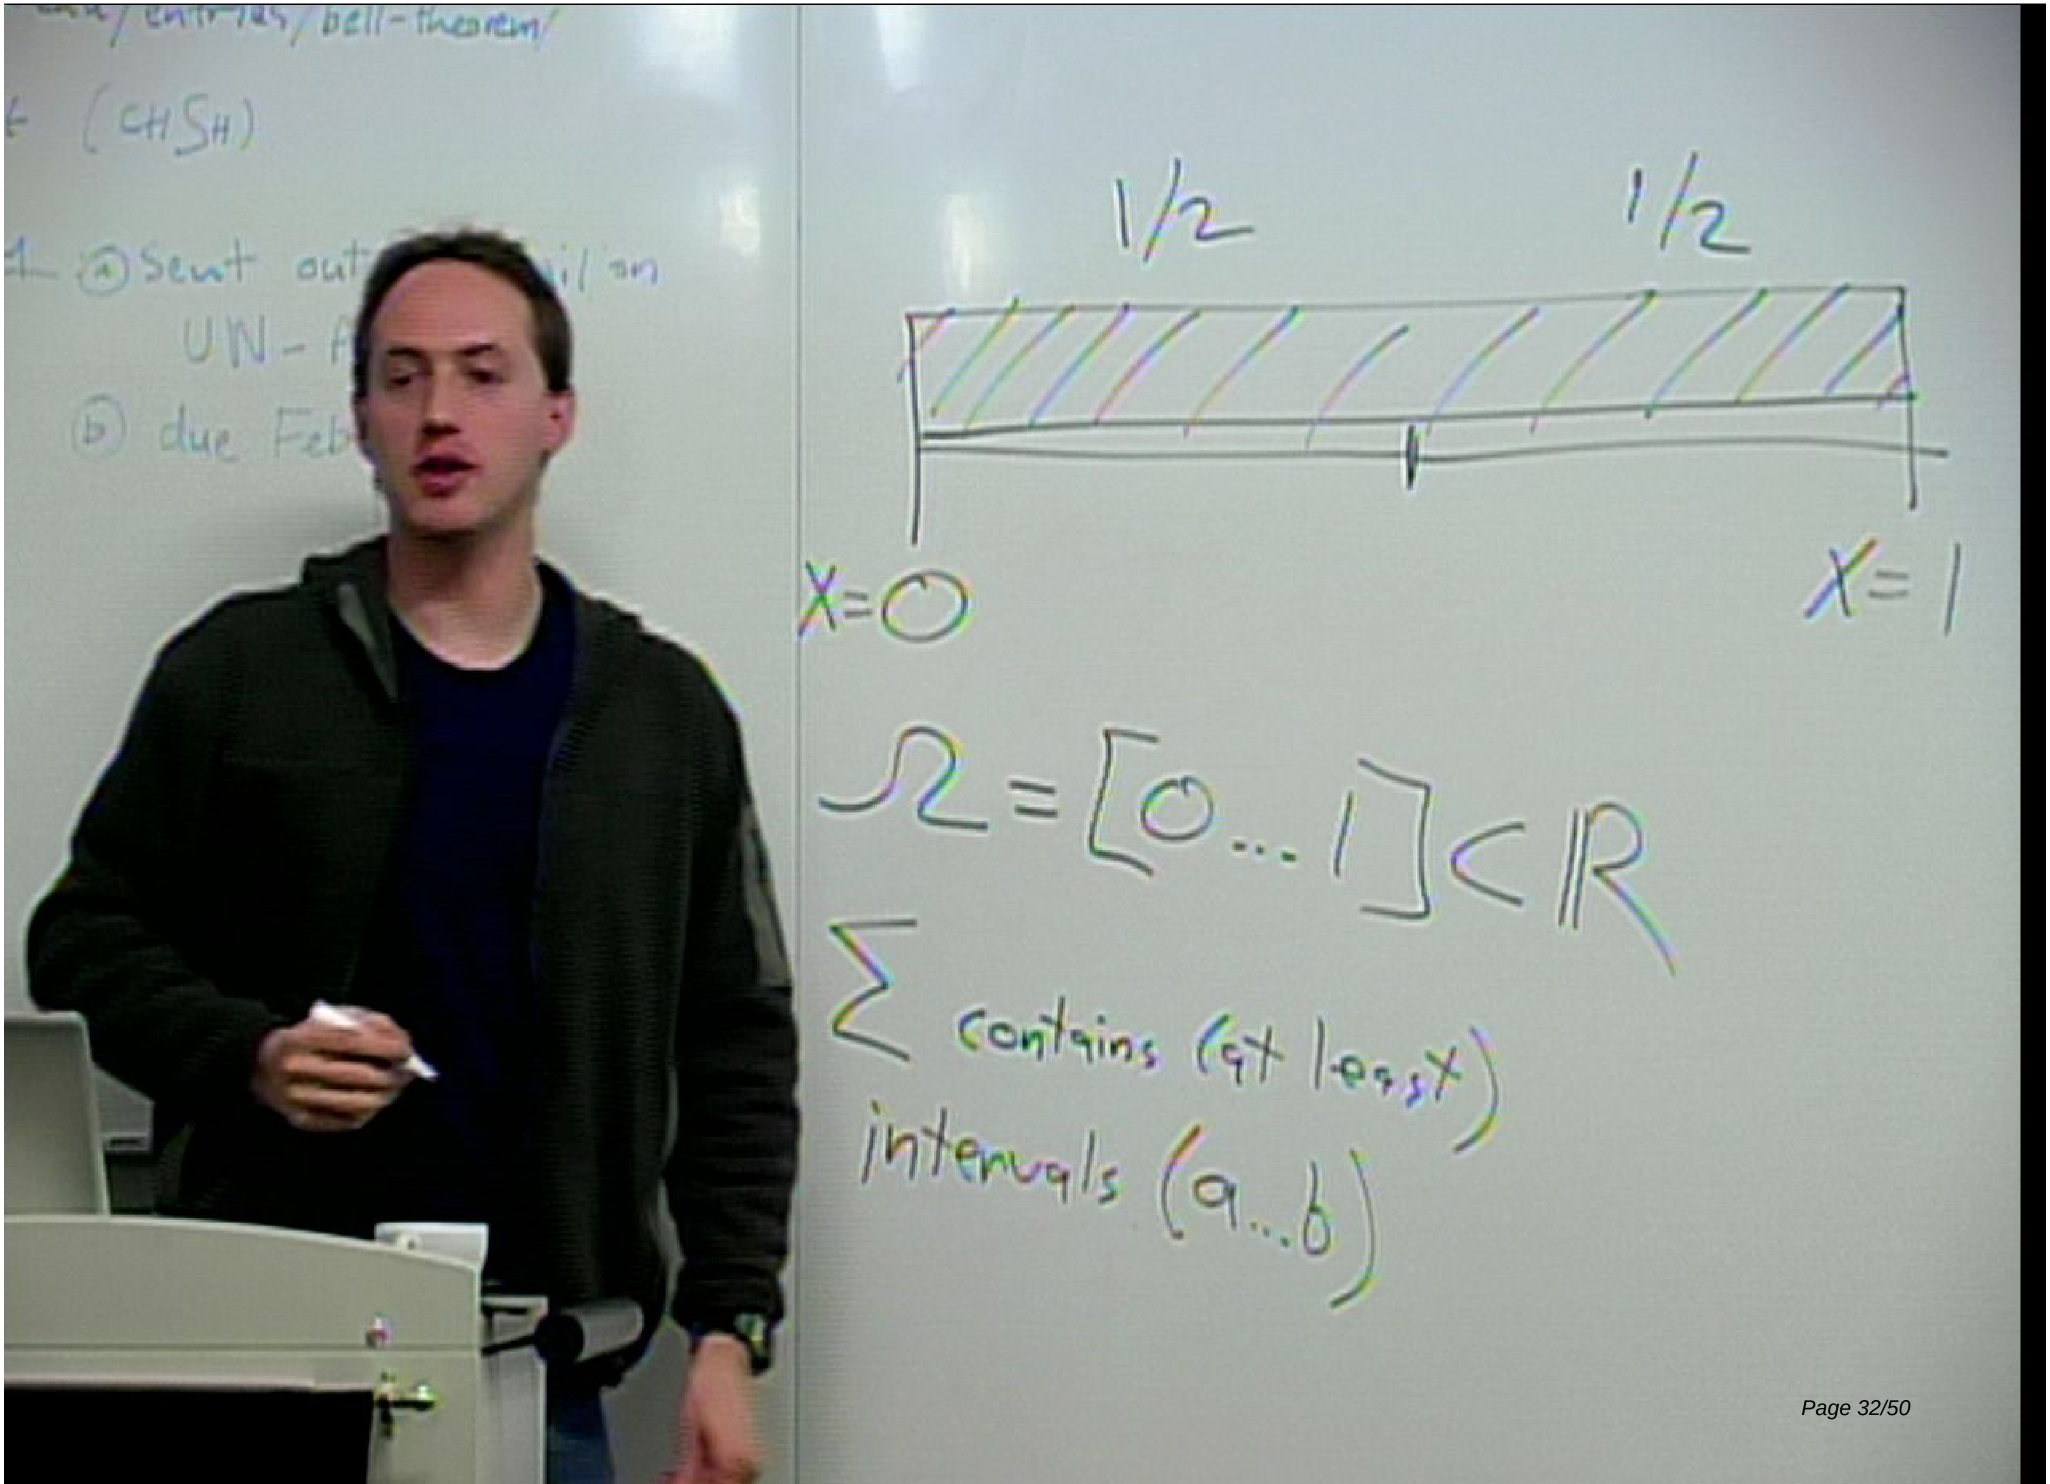
entries/bell-theorem/
(CHSH)
Sent out
UN-
due Feb
1/2
1/2
X=0
X=1
Ω=[0...1]⊂R
Σ contains (at least)
intervals (a...b)

All the rules of probability theory — in particular, its **coherence** — can be derived from these axioms. However, probability "is" nothing but a *measure* — i.e., a function that assigns real numbers to subsets of Ω. No connection to physics or real life is made. However, we can do calculations all day within this framework, interpreting the results however we please.

Note that Kolmogorov's axioms (a) are a bit more complex than you might think is necessary (e.g., "Can't we just assign probabilities to elementary events, and dispense with Σ?"), and (b) allow some surprising and counterintuitive possibilities. For instance:

1. Σ may not contain any events corresponding to single elementary events! This is why we can't just assign probabilities to elementary events.
2. If it does, it may be that every elementary event has $Pr(x) = 0$... if there are uncountably many elementary events.
3. $Pr(E) = 0$ does not imply that E cannot happen.
4. Some sample spaces have subsets that *cannot* appear in Σ — i.e., are "non-measurable".

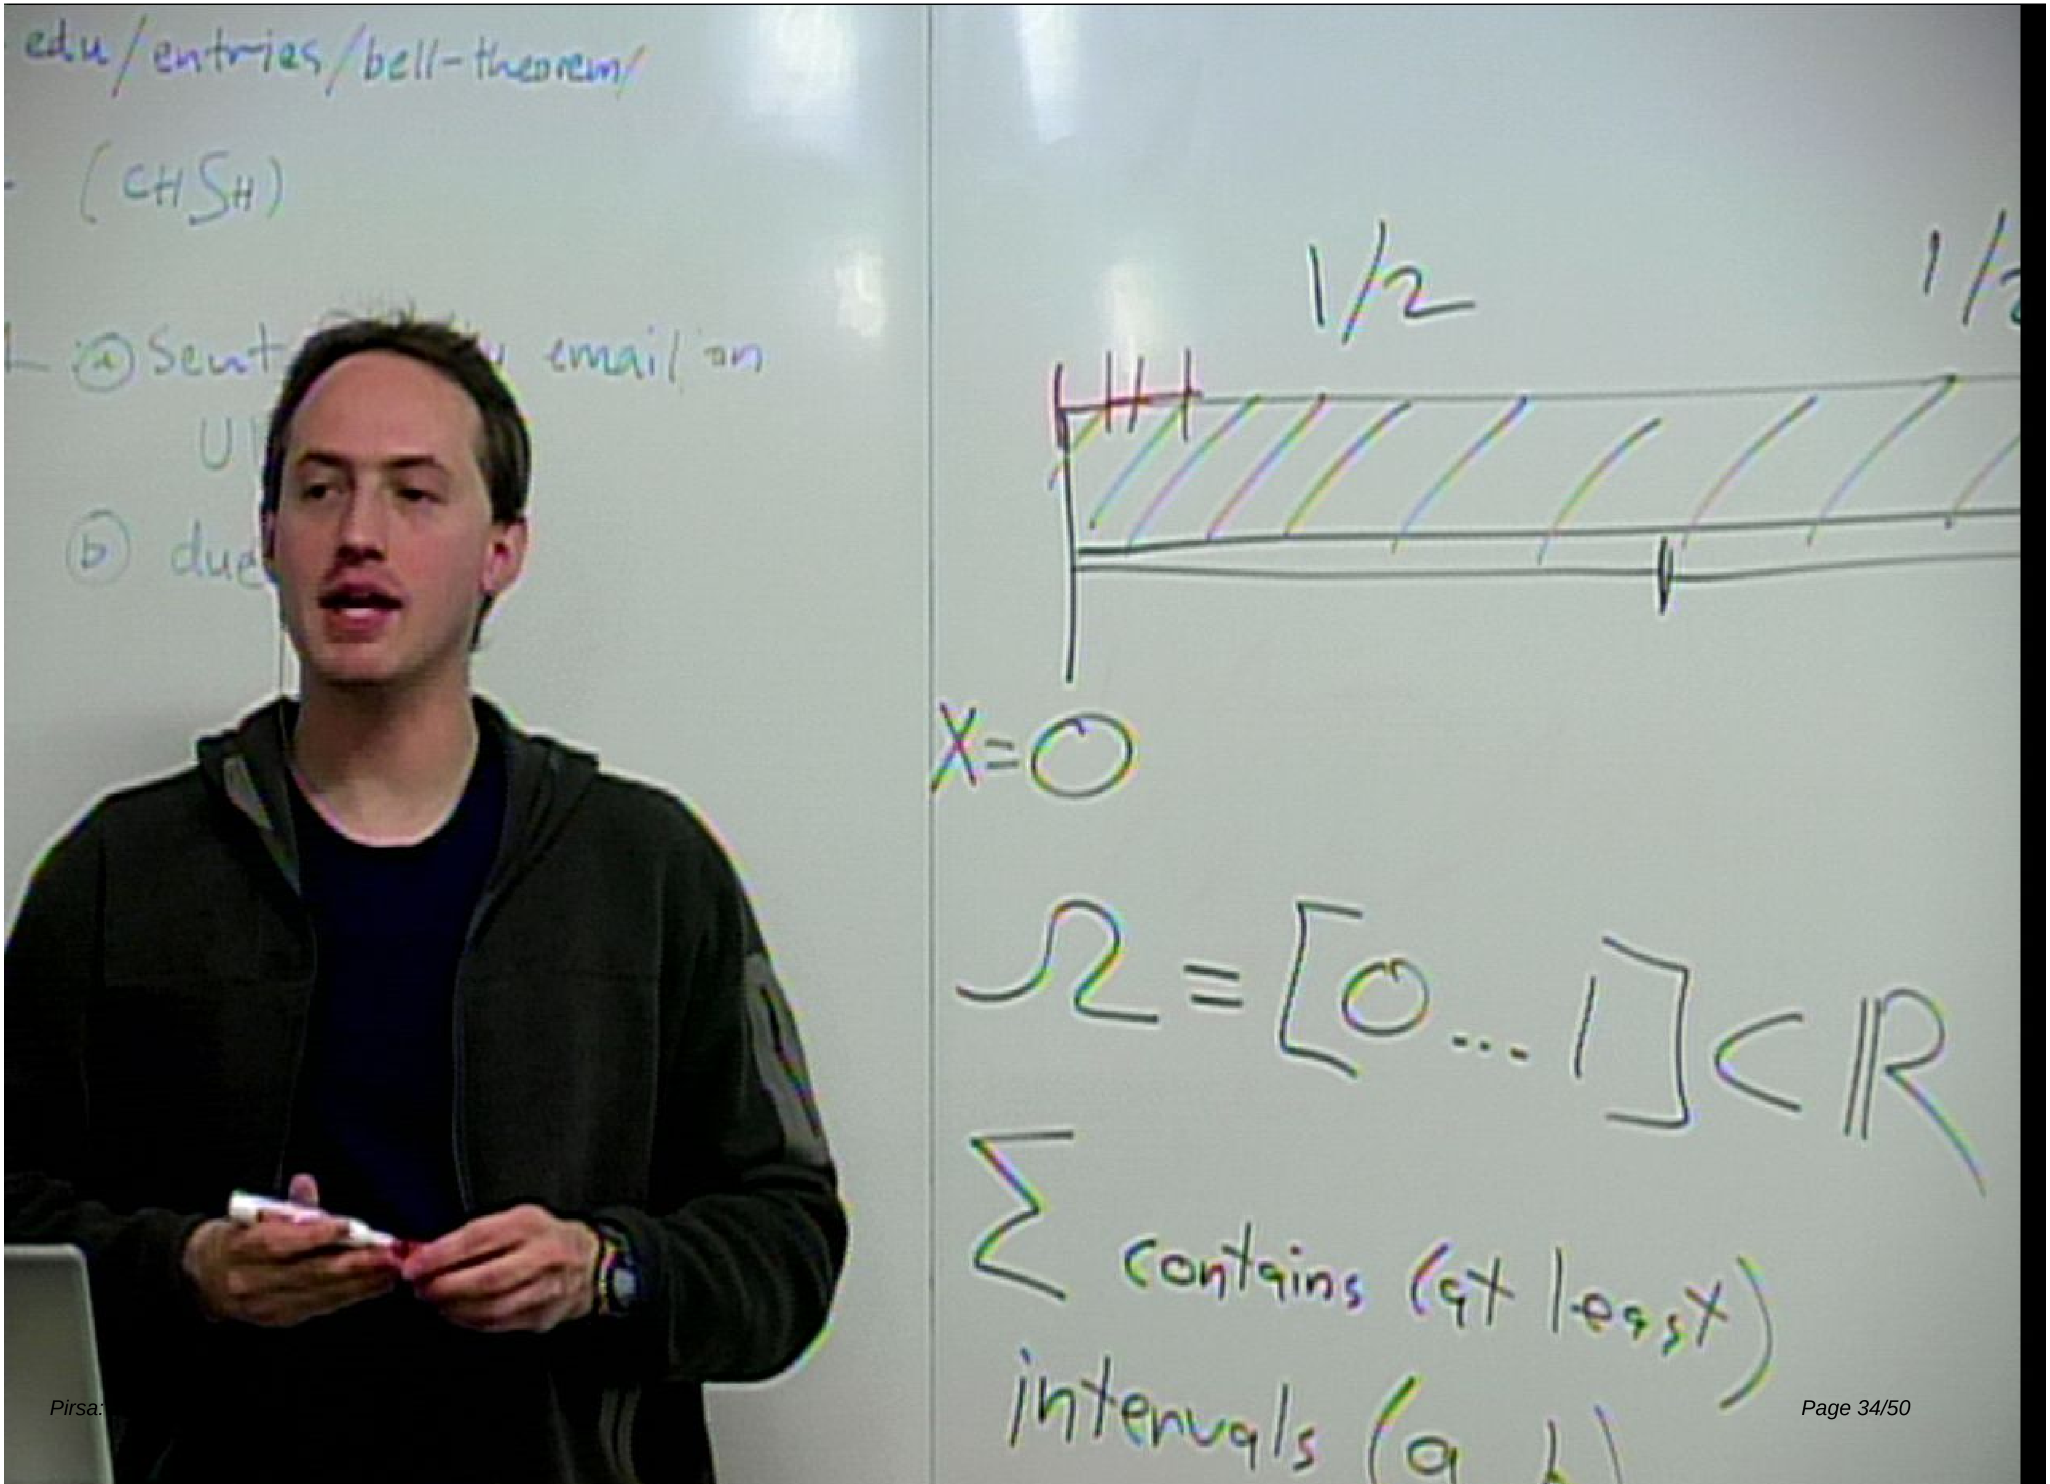
edu/entries/bell-theorem/
(CHSH)
1/2
X=0
Ω ≡ [0 ... 1] ⊂ R
Σ contains (at least)
intervals

Firefox File Edit View History Bookmarks Tools Window Help Thu 3:12 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

function that assigns real numbers to subsets of Ω. No connection to physics or real life is made. However, we can do calculations all day within this framework, interpreting the results however we please.

Note that Kolmogorov's axioms (a) are a bit more complex than you might think is necessary (e.g., "Can't we just assign probabilities to elementary events, and dispense with Σ?"), and (b) allow some surprising and counterintuitive possibilities. For instance:

1. Σ may not contain any events corresponding to single elementary events! This is why we can't just assign probabilities to elementary events.
2. If it does, it may be that every elementary event has $Pr(x) = 0$... if there are uncountably many elementary events.
3. $Pr(E) = 0$ does not imply that E cannot happen.
4. Some sample spaces have subsets that *cannot* appear in Σ — i.e., are "non-measurable".

Bayesians, by the way, have mechanisms to derive Kolmogorov's axioms from other axioms, which are considered plausible statements either about logic (Cox) or gambling (DeFinetti).

Firefox File Edit View History Bookmarks Tools Window Help Thu 3:12 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

from these axioms. However, probability "is" nothing but a *measure* — i.e., a function that assigns real numbers to subsets of Ω. No connection to physics or real life is made. However, we can do calculations all day within this framework, interpreting the results however we please.

Note that Kolmogorov's axioms (a) are a bit more complex than you might think is necessary (e.g., "Can't we just assign probabilities to elementary events, and dispense with Σ?"), and (b) allow some surprising and counterintuitive possibilities. For instance:

1. Σ may not contain any events corresponding to single elementary events! This is why we can't just assign probabilities to elementary events.
2. If it does, it may be that every elementary event has $Pr(x) = 0$... if there are uncountably many elementary events.
3. $Pr(E) = 0$ does not imply that E cannot happen.
4. Some sample spaces have subsets that *cannot* appear in Σ — i.e., are "non-measurable".

Bayesians, by the way, have mechanisms to derive Kolmogorov's axioms from other axioms, which are considered plausible statements either about logic (Cox) or gambling (DeFinetti).

function that assigns real numbers to subsets of Ω. No connection to physics or real life is made. However, we can do calculations all day within this framework, interpreting the results however we please.

Note that Kolmogorov's axioms (a) are a bit more complex than you might think is necessary (e.g., "Can't we just assign probabilities to elementary events, and dispense with Σ?"), and (b) allow some surprising and counterintuitive possibilities. For instance:

1. Σ may not contain any events corresponding to single elementary events! This is why we can't just assign probabilities to elementary events.
2. If it does, it may be that every elementary event has $Pr(x) = 0$... if there are uncountably many elementary events.
3. $Pr(E) = 0$ does not imply that E cannot happen.
4. Some sample spaces have subsets that *cannot* appear in Σ — i.e., are "non-measurable".

Bayesians, by the way, have mechanisms to derive Kolmogorov's axioms from other axioms, which are considered plausible statements either about logic (Cox) or gambling (DeFinetti).

Note that Kolmogorov's axioms (a) are a bit more complex than you might think is necessary (e.g., "Can't we just assign probabilities to elementary events, and dispense with Σ?"), and (b) allow some surprising and counterintuitive possibilities. For instance:

1. Σ may not contain any events corresponding to single elementary events! This is why we can't just assign probabilities to elementary events.
2. If it does, it may be that every elementary event has $Pr(x) = 0$... if there are uncountably many elementary events.
3. $Pr(E) = 0$ does not imply that E cannot happen.
4. Some sample spaces have subsets that *cannot* appear in Σ — i.e., are "non-measurable".

Bayesians, by the way, have mechanisms to derive Kolmogorov's axioms from other axioms, which are considered plausible statements either about logic (Cox) or gambling (DeFinetti).

What is the *use* of a probability?

What is the *use* of a probability?

Now that we have the machinery, let's play with it! The goal here is to get some idea what frequentists and Bayesians are thinking when they think about concrete problems. *Gambling* is the oldest problem in the book, and while it's the backbone of Bayesian theory, everybody does it (even frequentists!) A more modern problem that may be interesting to those of you from CS or information theory is *data compression*.

One of the virtues of frequentism is that a frequentist can very easily say "Here is a machine that flips coins, and each coin that is flipped has probabilities $\vec{P} = (p, 1-p)$ for landing heads or tail, respectively." Or just "Consider a sequence of coin flips with distribution $\vec{P} = (p, 1-p)$.

Saying the same thing in Bayesian is like writing "Hello, World" in assembly language. But Bayesians eventually learned how to say it, as "For every $N > 0$, I have a probability distribution over the 2^N-element set $\{heads, tails\}^N$, which is given by $\vec{P}^{\otimes N}$."

Firefox File Edit View History Bookmarks Tools Window Help Thu 3:14 PM Robin Blume-Kohout

Robin's Notes – An ongoing quest for coherence

Robin's Notes – An o... | Bayesian probability ... | Probability axioms – ... | Quantum Theory – | Principia Mathematic...

What is the *use* of a probability?

Now that we have the machinery, let's play with it! The goal here is to get some idea what frequentists and Bayesians are thinking when they think about concrete problems. *Gambling* is the oldest problem in the book, and while it's the backbone of Bayesian theory, everybody does it (even frequentists!) A more modern problem that may be interesting to those of you from CS or information theory is *data compression*.

One of the virtues of frequentism is that a frequentist can very easily say "Here is a machine that flips coins, and each coin that is flipped has probabilities $\vec{P} = (p, 1-p)$ for landing heads or tail, respectively." Or just "Consider a sequence of coin flips with distribution $\vec{P} = (p, 1-p)$.

Saying the same thing in Bayesian is like writing "Hello, World" in assembly language. But Bayesians eventually learned how to say it, as "For every $N > 0$, I have a probability distribution over the 2^N-element set $\{heads, tails\}^N$, which is given by $\vec{P}^{\otimes N}$."

Gambling

Okay, let's consider a sequence of coin flips with distribution $\vec{P} = (p, 1-p)$. Let's play a game.

1. You have D dollars.
2. Before the coin is flipped, you divide your money into two piles, one labeled "It will be heads" and one labeled "It will be tails".
3. The coin is flipped.
4. I claim the money from the false pile, but I *double* the money in the true pile.

So, suppose you divide your money in proportions $x, 1-x$. Then, after the flip:
1. with probability p, you have $2x$ times your original money,
2. with probability $1-p$, you have $2(1-x)$ times your original money.

Note that the game is explicitly scale-invariant, so you can play it many times in a row. How should you play?

You might start by maximizing your *expected* post-game money. This is

$$\langle D\rangle_1 = [2xp + 2(1-x)(1-p)]D = 2D[(1-p) + (2p-1)x].$$

Okay, let's consider a sequence of coin flips with distribution $\vec{P} = (p, 1-p)$. Let's play a game.

1. You have D dollars.
2. Before the coin is flipped, you divide your money into two piles, one labeled "It will be heads" and one labeled "It will be tails".
3. The coin is flipped.
4. I claim the money from the false pile, but I *double* the money in the true pile.

So, suppose you divide your money in proportions $x, 1-x$. Then, after the flip:
1. with probability p, you have $2x$ times your original money,
2. with probability $1-p$, you have $2(1-x)$ times your original money.

Note that the game is explicitly scale-invariant, so you can play it many times in a row. How should you play?

You might start by maximizing your *expected* post-game money. This is

$$\langle D \rangle_1 = [2xp + 2(1-x)(1-p)]D = 2D[(1-p) + (2p-1)x].$$

If $2p-1$ is positive — i.e., $p > \frac{1}{2}$ — then you should set $x = 1$ and bet *all* your

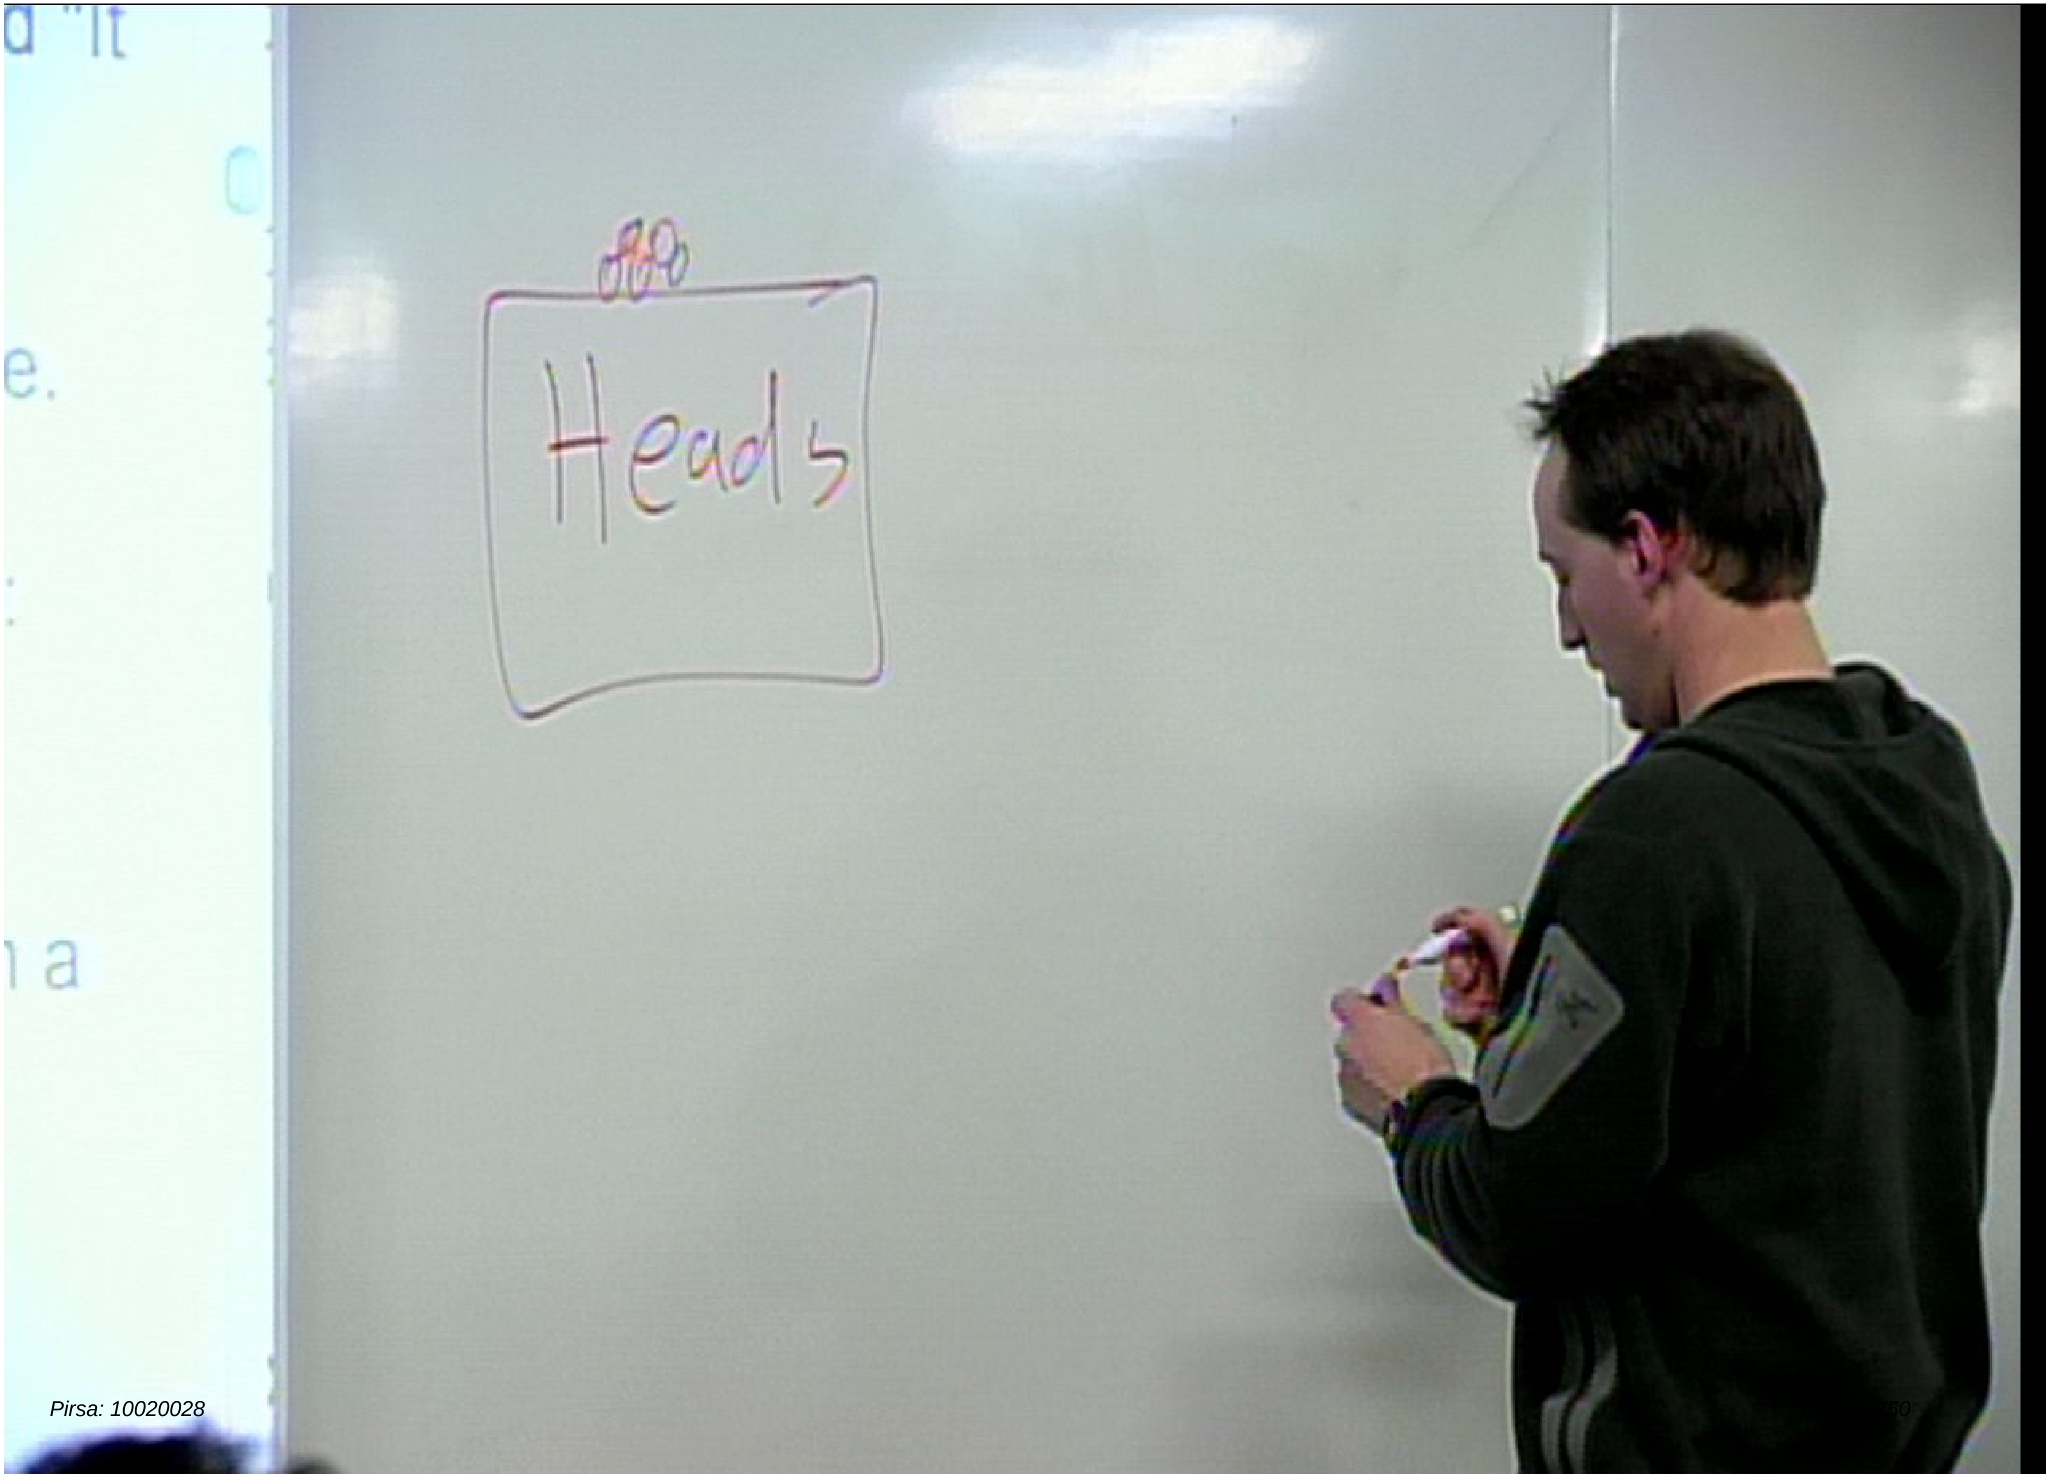
Heads

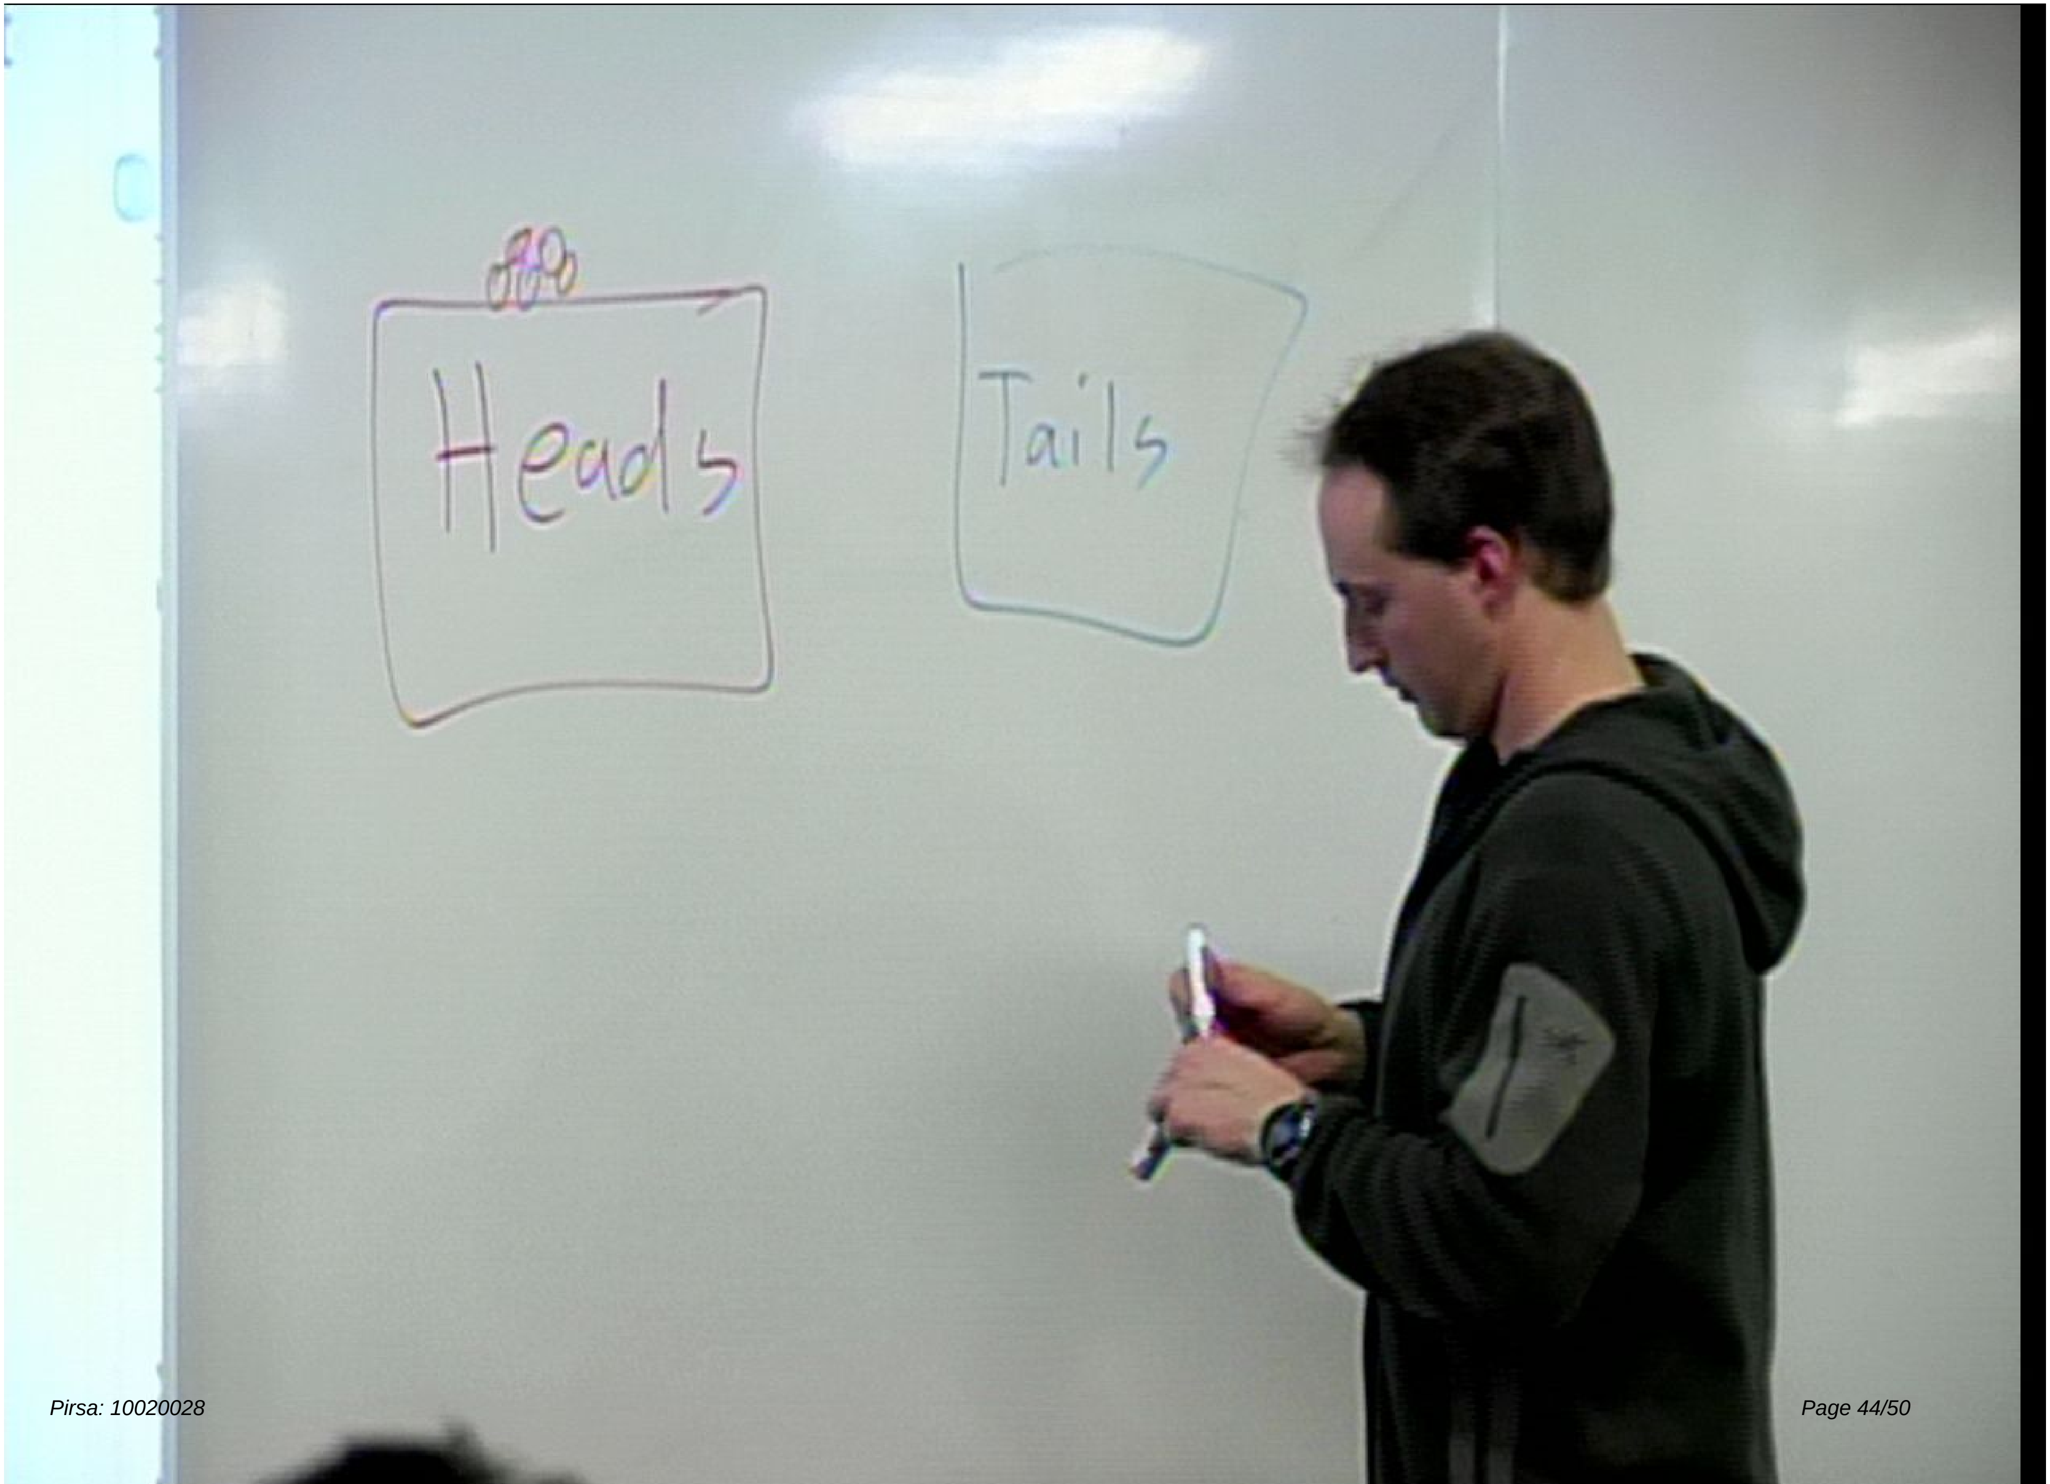
Heads
Tails

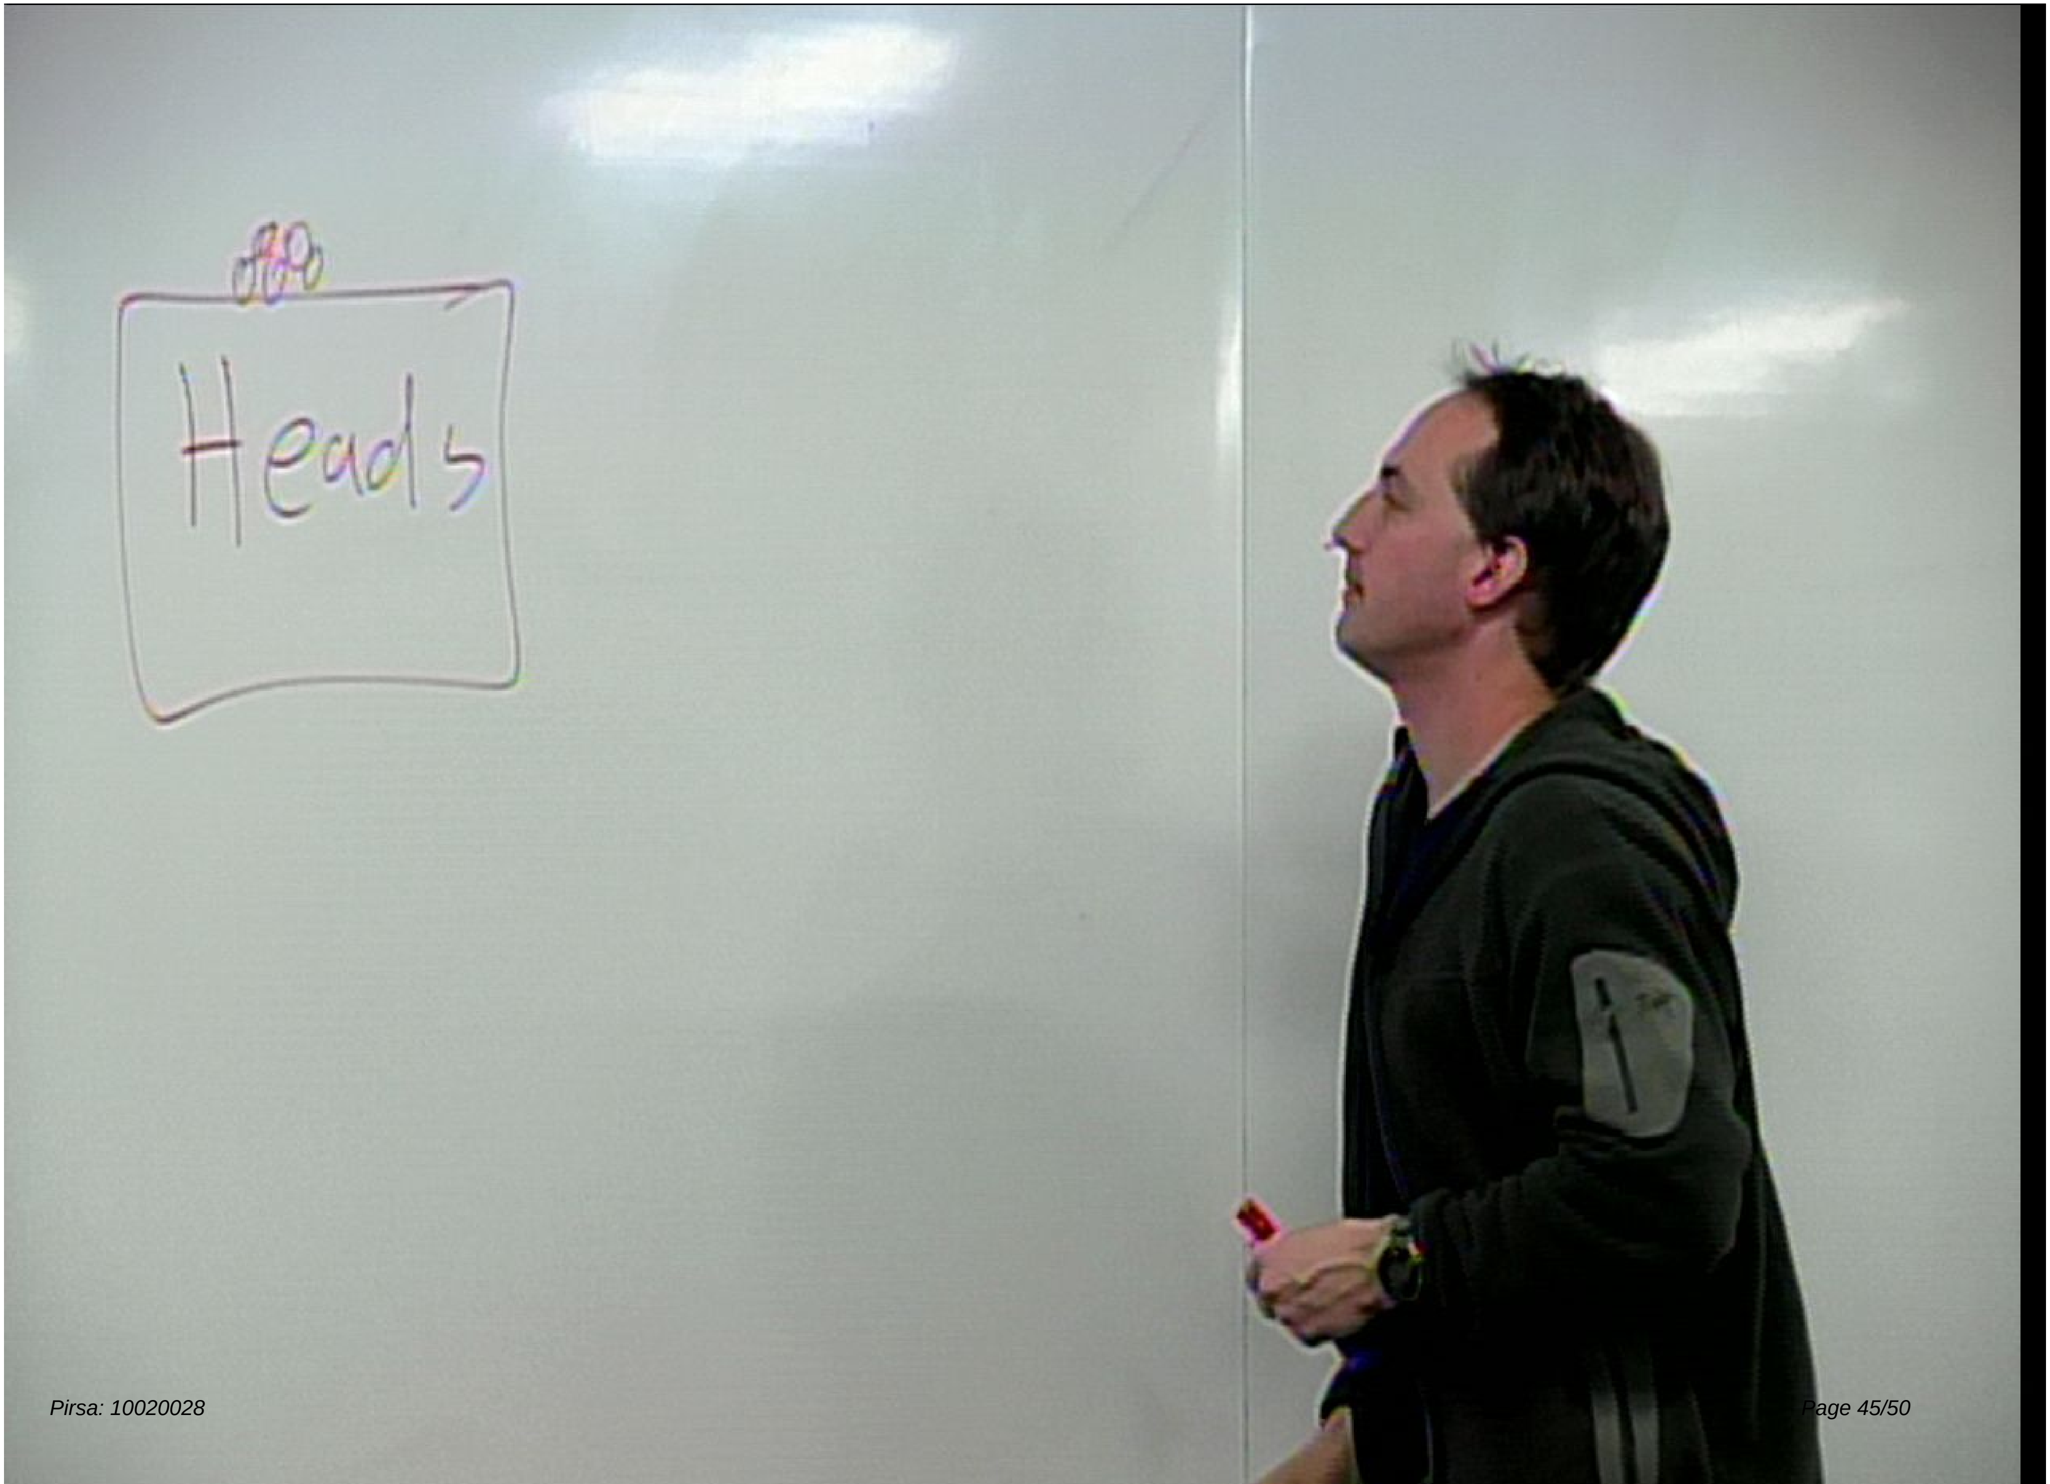
Heads

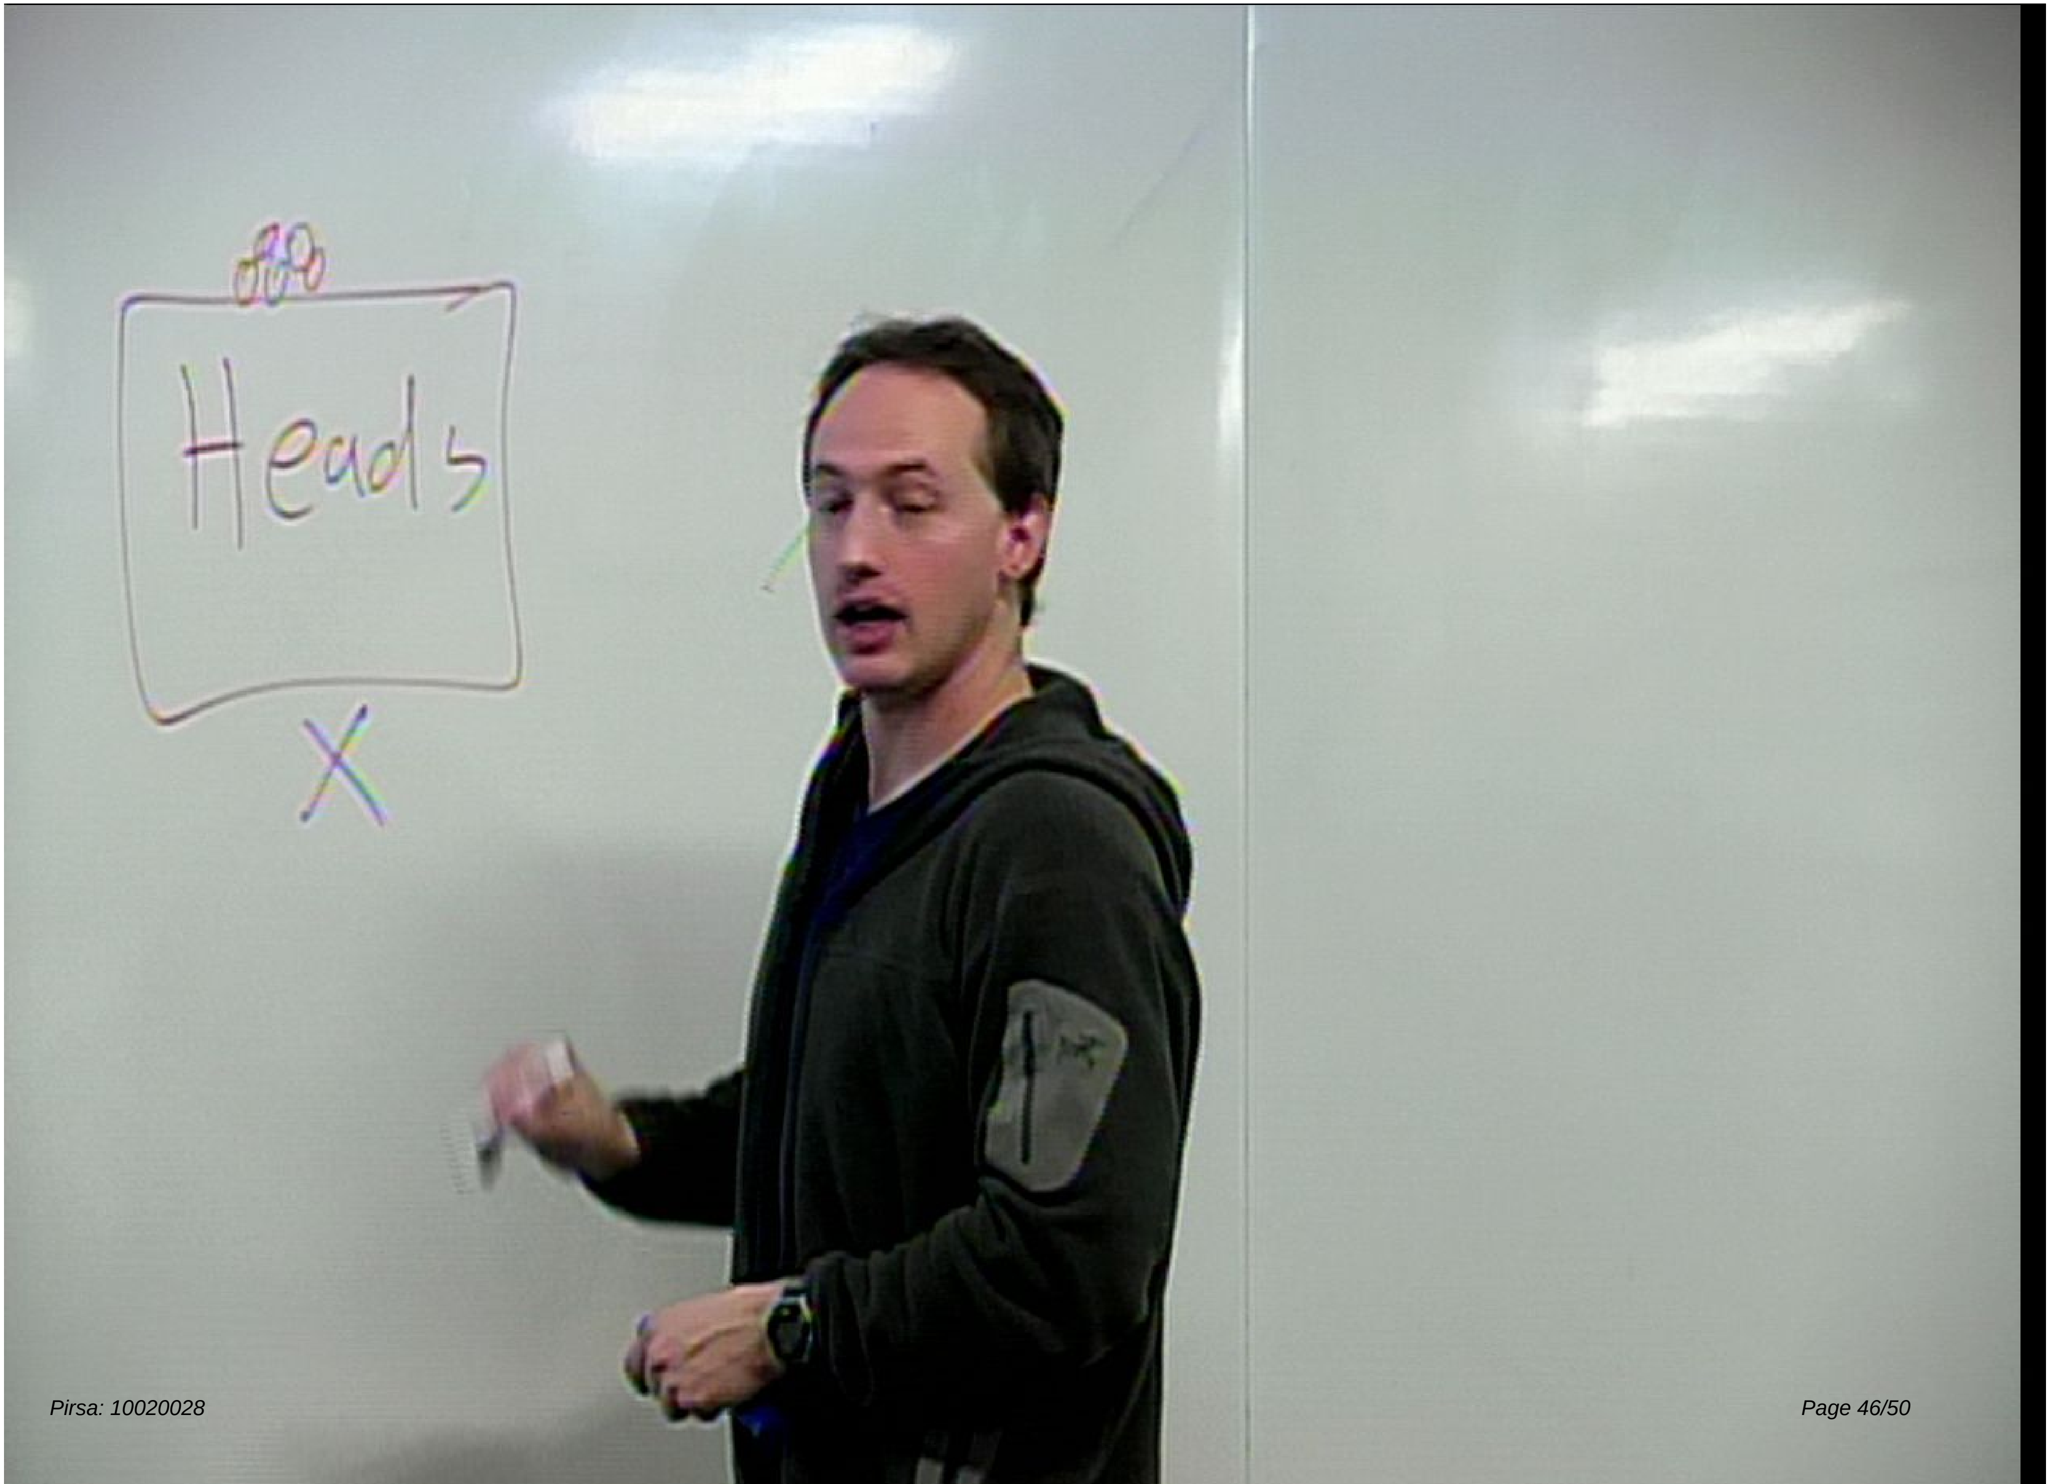
Heads
X

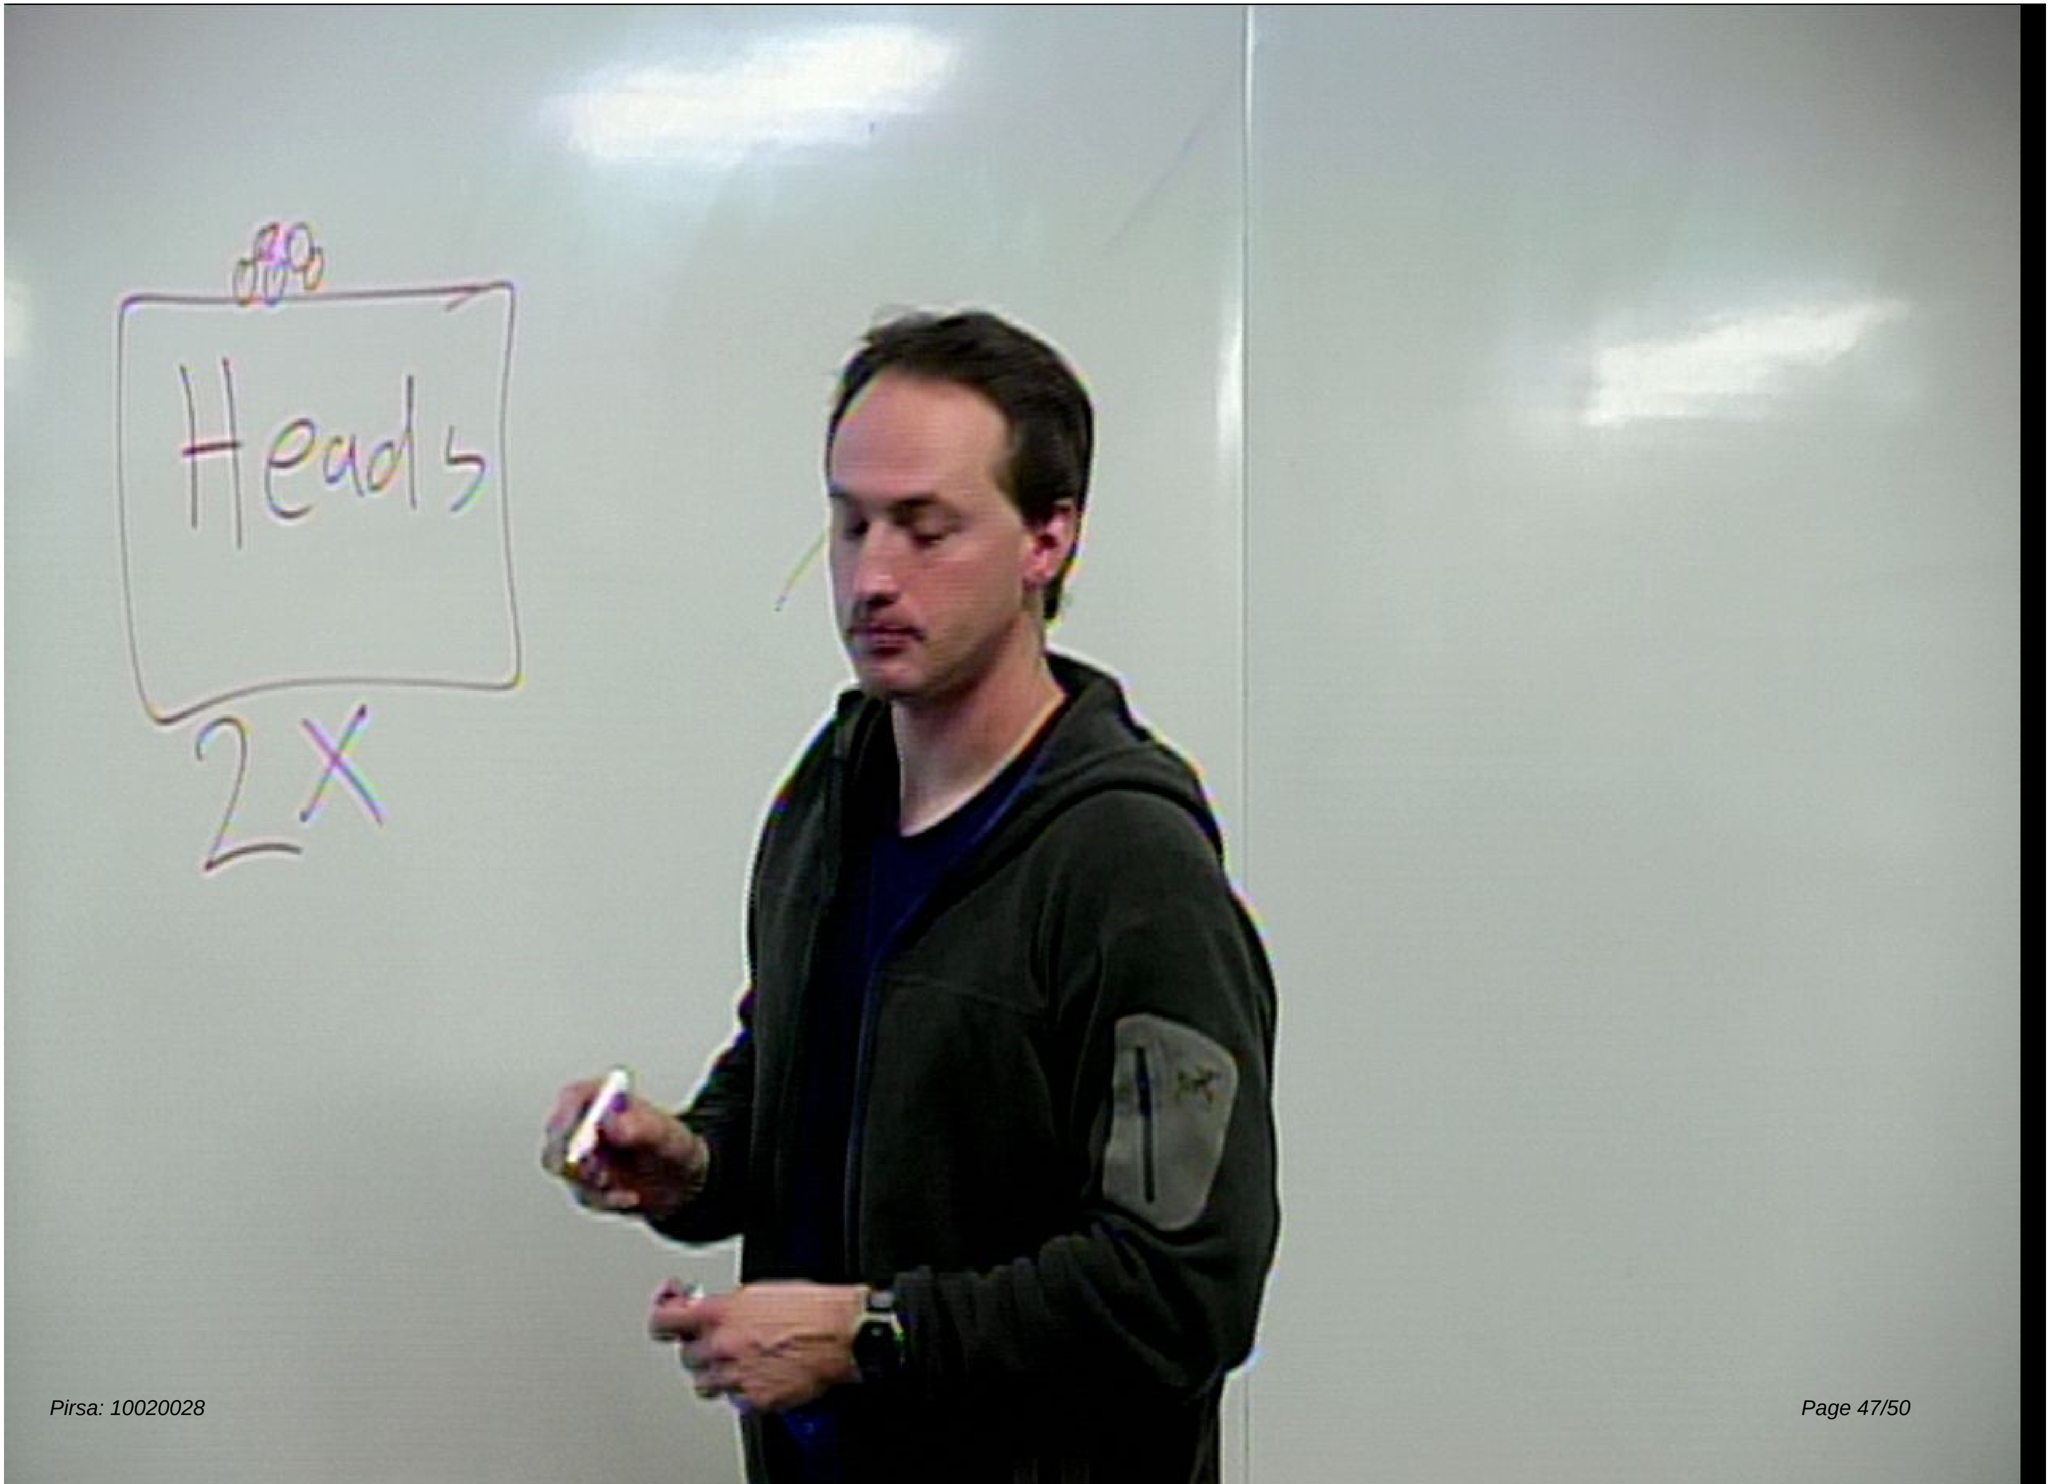
Heads
2X

Firefox File Edit View History Bookmarks Tools Window Help Thu 3:19 PM Robin Blume-Kohout

Robin's Notes - An ongoing quest for coherence

Robin's Notes - An o... | Bayesian probability ... | Probability axioms - ... | Quantum Theory - | Principia Mathematic...

Okay, let's consider a sequence of coin flips with distribution $\vec{P} = (p, 1-p)$. Let's play a game.

1. You have D dollars.
2. Before the coin is flipped, you divide your money into two piles, one labeled "It will be heads" and one labeled "It will be tails".
3. The coin is flipped.
4. I claim the money from the false pile, but I *double* the money in the true pile.

So, suppose you divide your money in proportions $x, 1-x$. Then, after the flip:
1. with probability p, you have $2x$ times your original money,
2. with probability $1-p$, you have $2(1-x)$ times your original money.

Note that the game is explicitly scale-invariant, so you can play it many times in a row. How should you play?

You might start by maximizing your *expected* post-game money. This is

$$\langle D \rangle_1 = [2xp + 2(1-x)(1-p)]D = 2D[(1-p) + (2p-1)x].$$

If $2p-1$ is positive — i.e., $p > \frac{1}{2}$ — then you should set $x = 1$ and bet *all* your

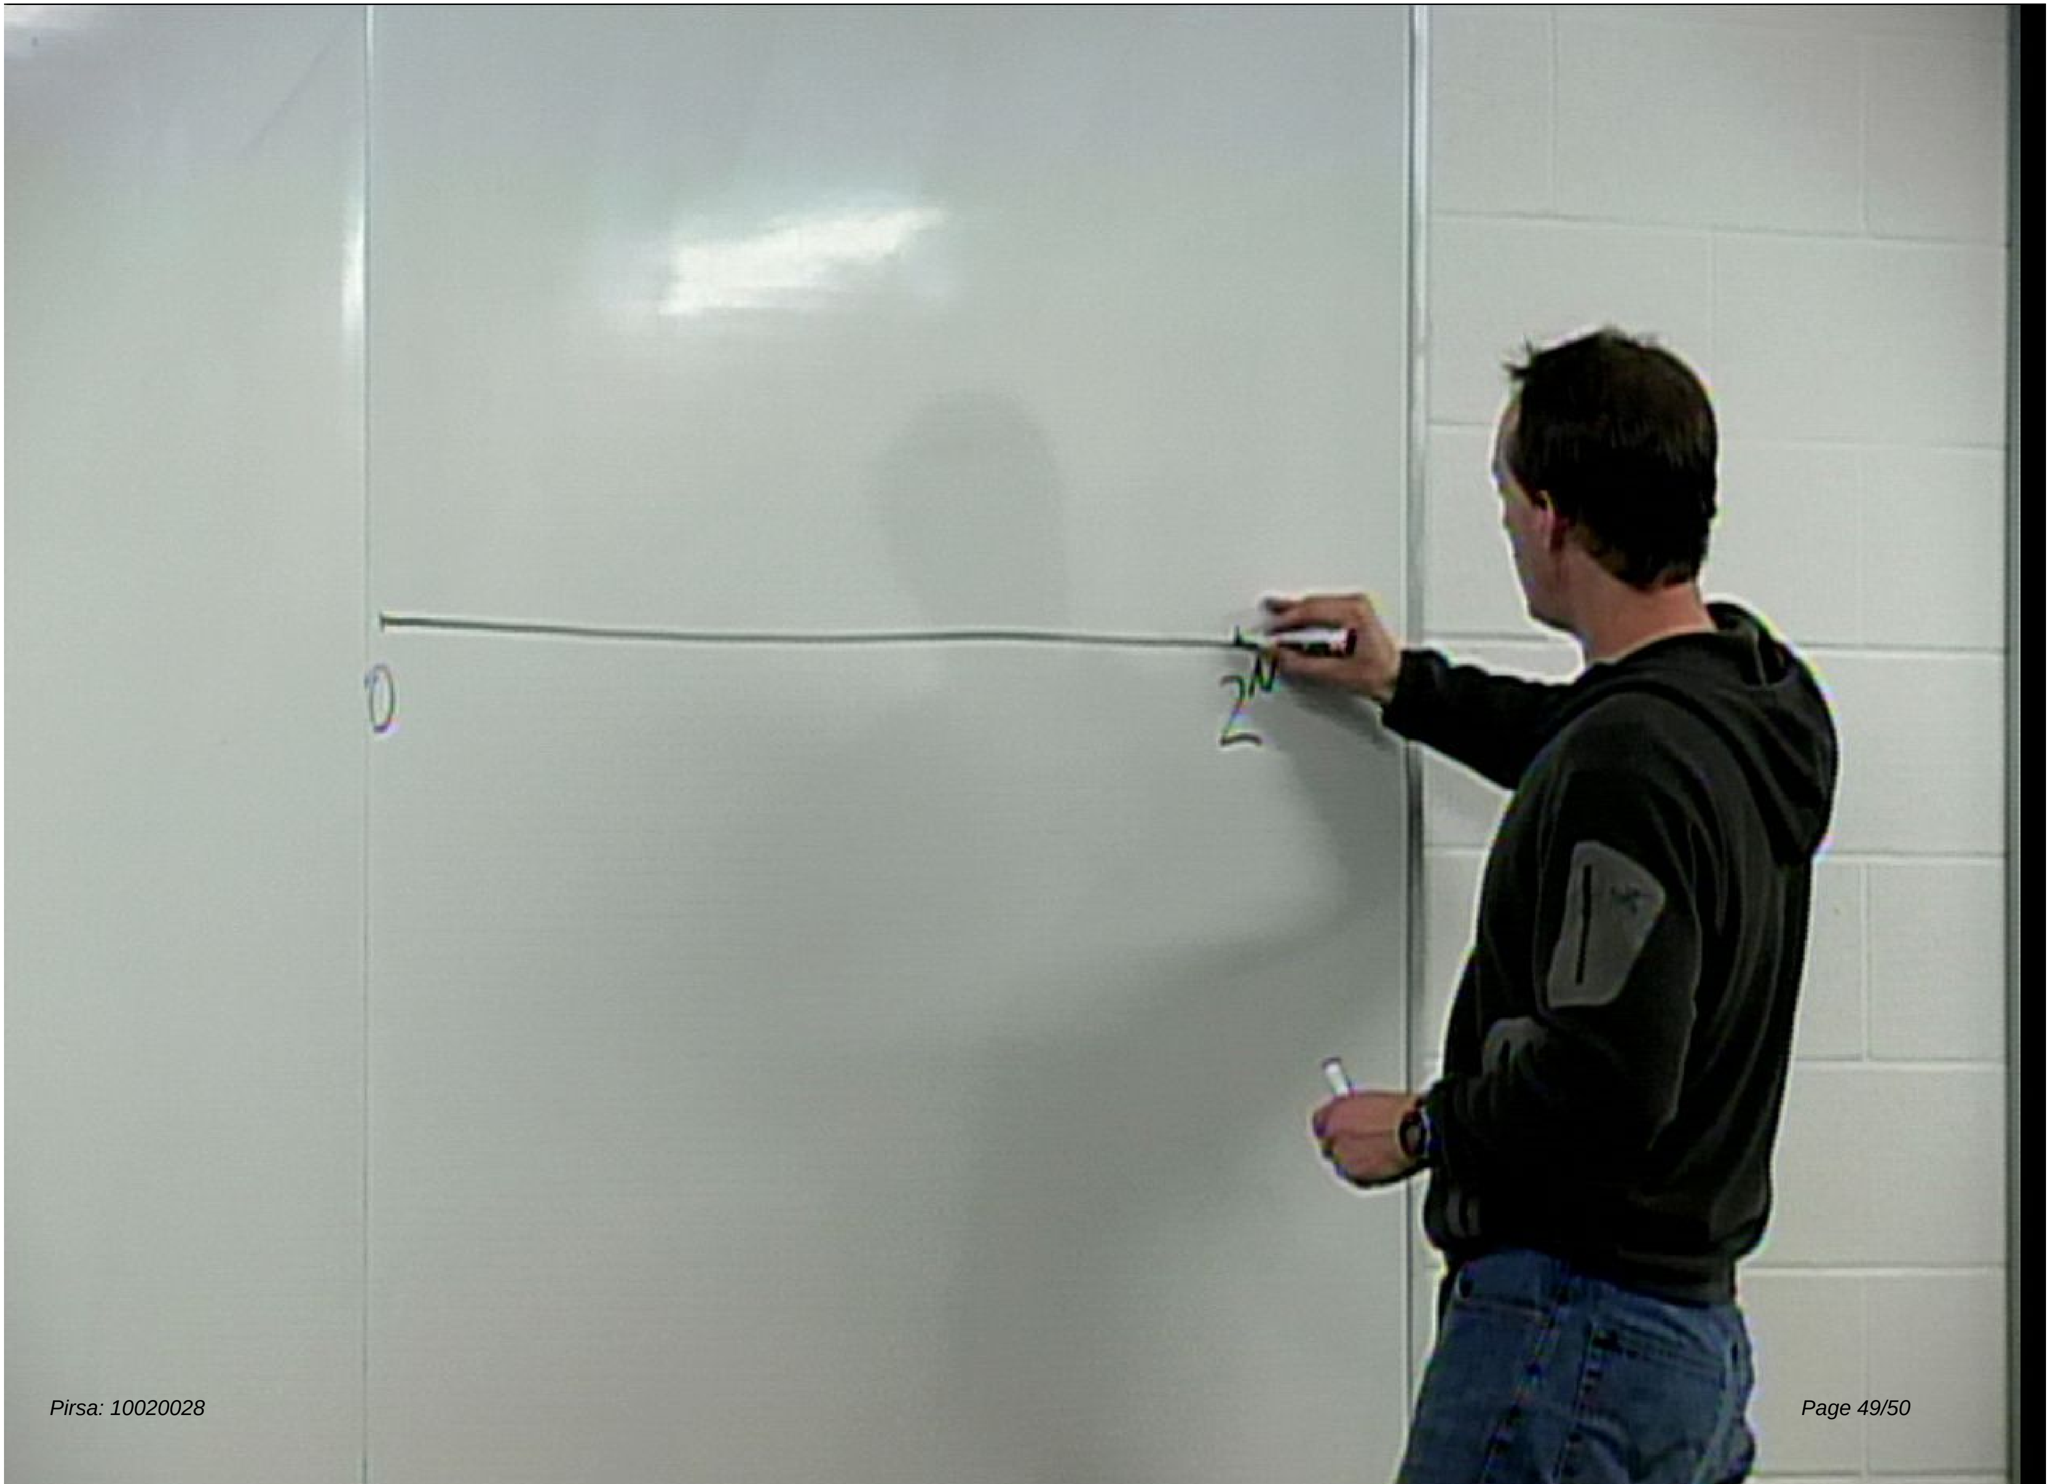
0
2^N

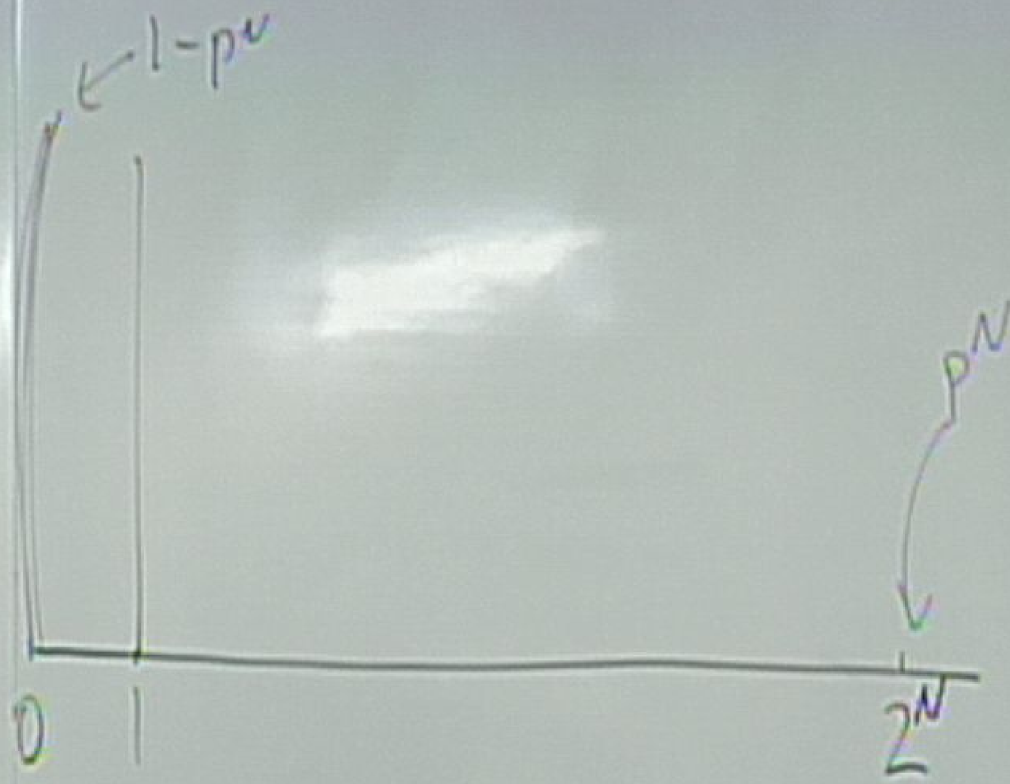
$1-p^N$
p^N
0
1
2^N