10/27/26: Macro finance seminar, Lars Peter Hansen

"Asset pricing, decision theory, and public policy design"

See Souganidis paper, ask Ethan about viscosity solutions, ask Anne about mean-field games.

"Intertemporal marginal social valuations" [SJ: ?] (exs: social cost of carbon, social cost of R&D)

Math trick: Girsanov

App: big project R&D motivated by climate change (climate economics, environmental economics, ask Tucker and AB)

Formulation is nontrivial; and "model interpretation", not "model solution"

"We confront broad-based uncertainties arising from imperfect measurement and quantification"

The NASEM released a 2017 recommendation on the social cost of carbon, with four distinct modules: the specialists subdivide, passing evaluations from one to the next. E.g., socioeconomic $\to$ climate $\to$ damage $\to$ discounting

But this is lacking an "integrated science" approach...

Game plan: build a value function; form "stochastic impulse responses"; deduce pricing formula; construct an uncertainty-averse tilted probability measure; model big changes as Poissonian (small shocks are modeled as Brownian increments); deduce "additive decompositions" of the contributions to the marginal valuations.

The state evolves according to:

$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t$ (drop the $(\cdot)$'s going forward)

Construct a "stochastic IR" to measure the impact of a "marginal change in a state"

We want to leverage the *theory of stochastic flows* by Kunita. [SJ: Kunita of Kunita–Watanabe fame? See my stochastic calculus notes.] Perturb the initial condition X_0. The stochastic IR process Λ captures the effect on future X's. [SJ: Formally, "$\Lambda := dX/dX_0$"?]

Component i of the IR, denoted Λ^i, evolves as

$d\Lambda_t^i = (\Lambda_t)'\partial_x\mu_i dt + (\Lambda_t)'\partial_x\sigma_i dW_t$ (basically, chain rule)

[SJ: Derive this for fun, learn stochastic flows.]

Stack all the Λ^i's and study joint dynamics of (Λ, X)

Linear VAR counterpart is a special case ($\mu(x) = cx, \sigma(x) \equiv \sigma$); in this case, Λ is deterministic. [SJ: Derive this for fun.]

[SJ: Q from audience: which components of X are endogenous vs exogenous? Reply: handle it later.]

Now we construct a value function associated with a disconuted objective with discount rate δ and utility U. There are maximizing controls (investments and emissions) and a minimizing probability measure.

Initial pricing formula: partial derivatives as asset prices

$\partial_x V(X_0) \cdot \Lambda_0 = \delta \int_0^\infty \exp(-\delta t) \mathbb{E}_0[\partial_x U(X_t) \cdot \Lambda_t] dt$

[SJ: How to read this equation? Summing the contributions of "marginal utility contributions in the future"; this is the term $\partial_x U(X_t) \cdot \Lambda_t$.]

[SJ: Q from Sean about changing coordinates to simplify Λ. Reply unclear to me.]

How does this differ from a co-state formulation of optimal control? No optimization ("envelope condition"), vector of interconnected co-state equations to solve forward, etc... [SJ: ?] This is easier to interpret.

State-dependent policies: $\mu(x) = \hat{\mu}[x, \Gamma(x)]$ and $U(x) = \hat{U}[x, \Gamma(x)]$, where Γ is the "policy rule". (Here, we're just making explicit the dependence of μ, U on the policy rule.)

The marginal values are $\delta \partial_\gamma \hat{U} \cdot \Delta$ and $\partial_x V \cdot [\partial_{\gamma^\top} \hat{\mu} \Delta]$. [SJ: ?]

Adding in jumps: imagine there are L kinds of jumps, with different Poisson process intensities. Let V^ℓ denote the "continuation value function" for the first jump being of type ℓ, conditional expectation.

Jump contributions. [SJ: Add picture. It would've been helpful for him to label each term in these chain rule equations. As Sean clarified, this term from time t (which itself is a sum over possible jumps $\ell \in [L]$) gets discounted back to time zero via $\delta \exp(-\delta t)$, then integrated over time $t \in [0, +\infty)$, as in the marginal value formula before.]

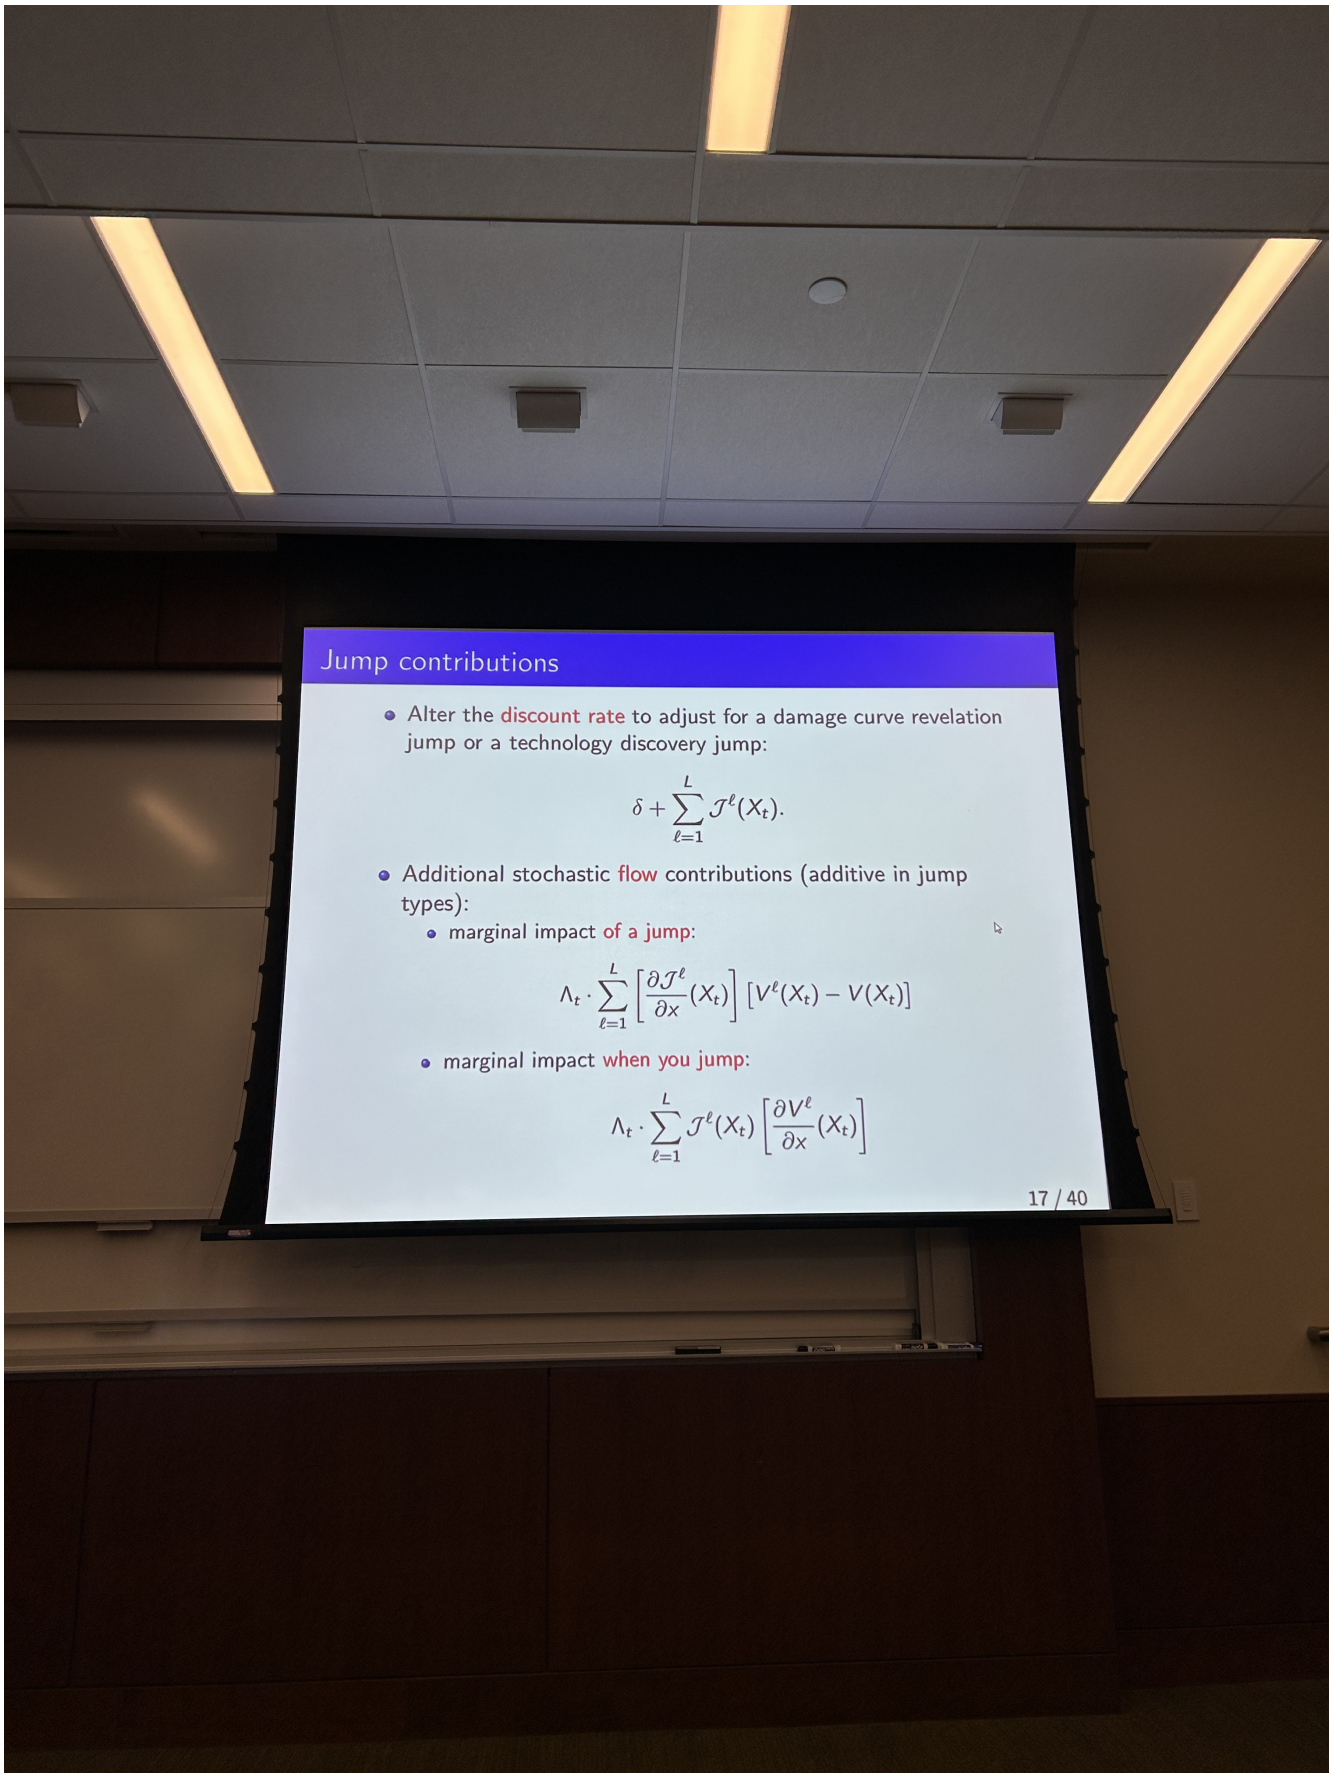


Figure 3: Formula for marginal contributions of jumps at time t.

Part II: adjusting valuation for uncertainty.

The term "deep uncertainty" shows up in the literature, due to the complexity of climate change forecasting. Borrow insights from decision theory.

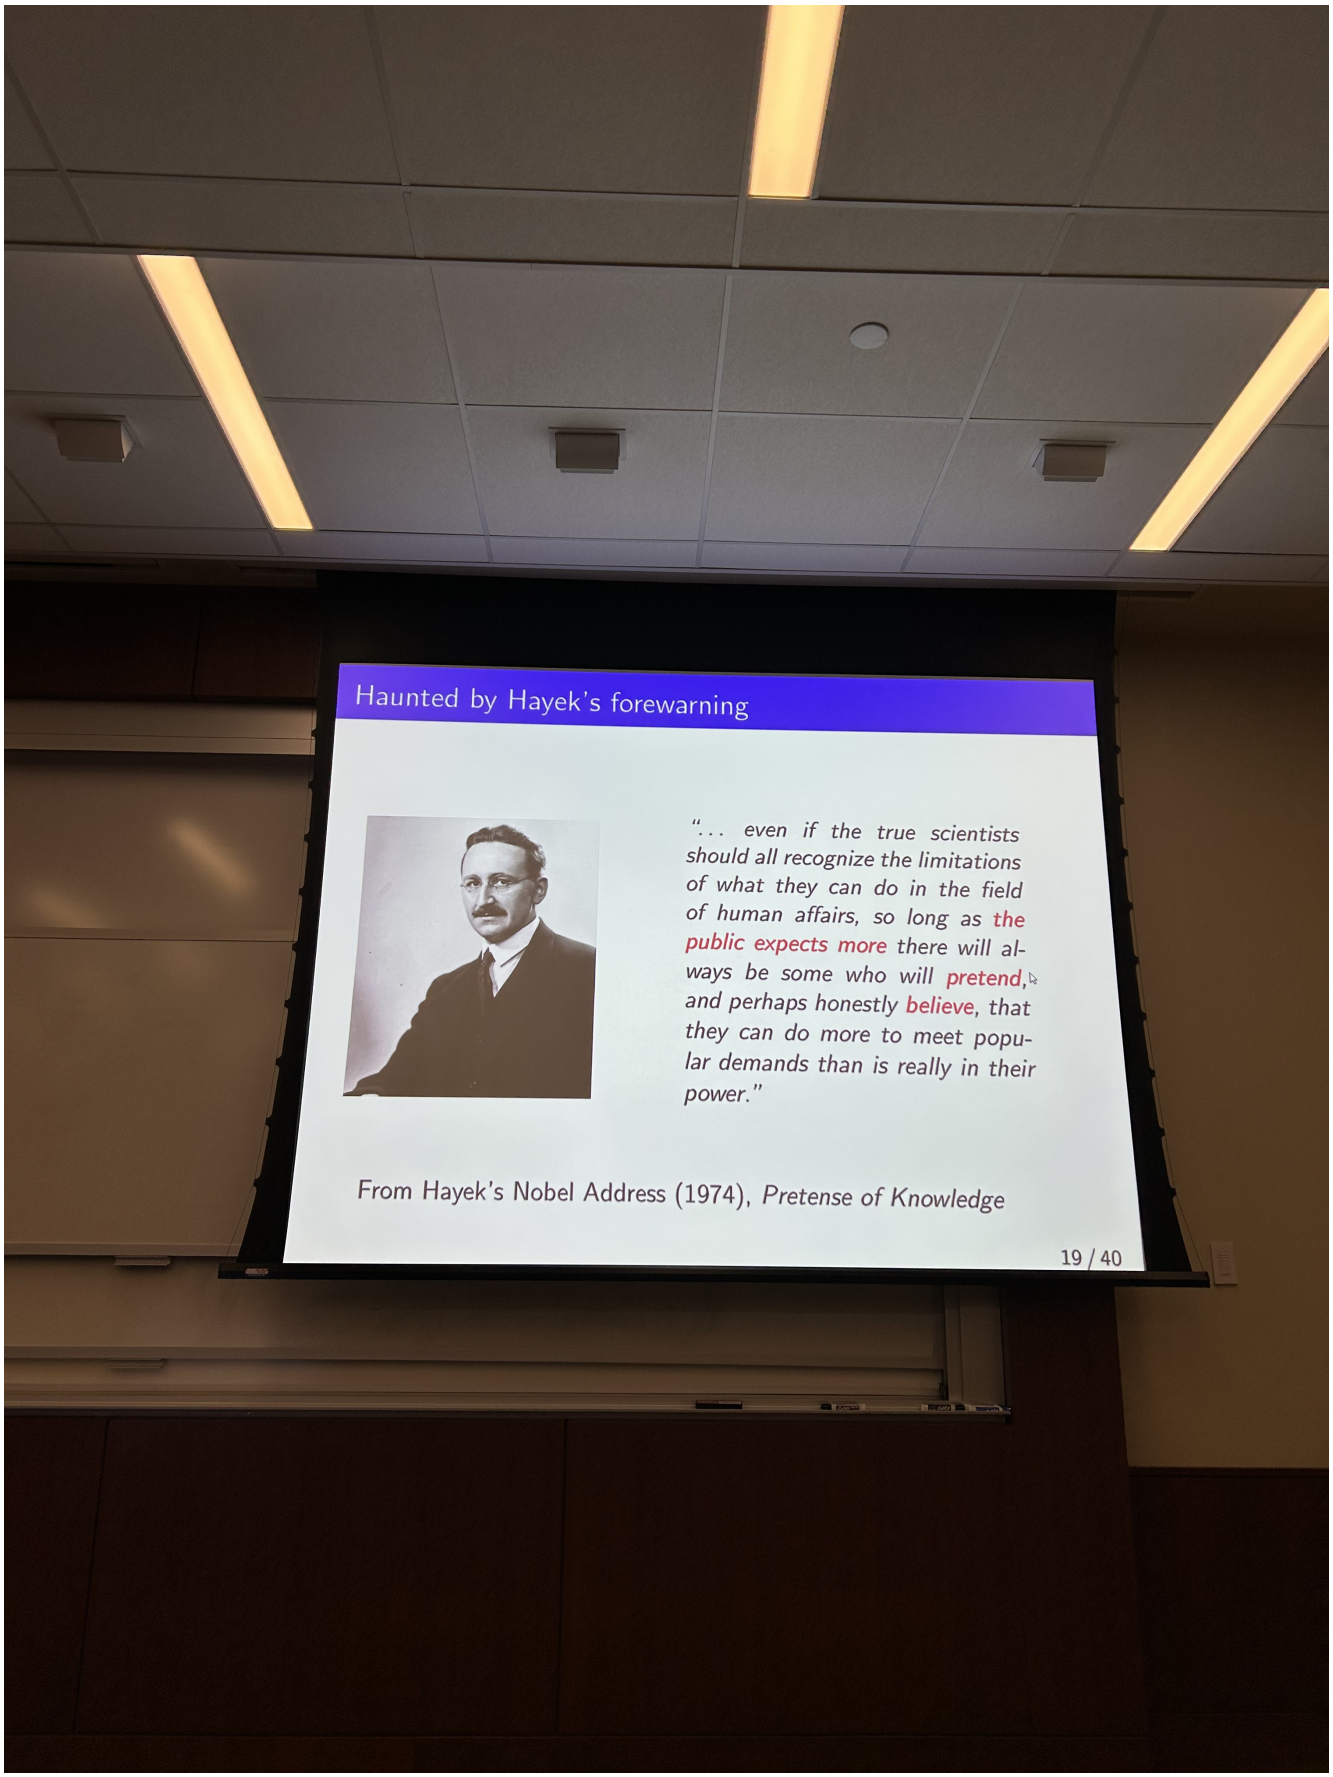


Figure 4: Is economics even a science? In this essay, Hayek roasts econometrics.

Economists, when consulted by policymakers, are often afraid to express uncertainty, because some policymakers will then question their legitimacy.

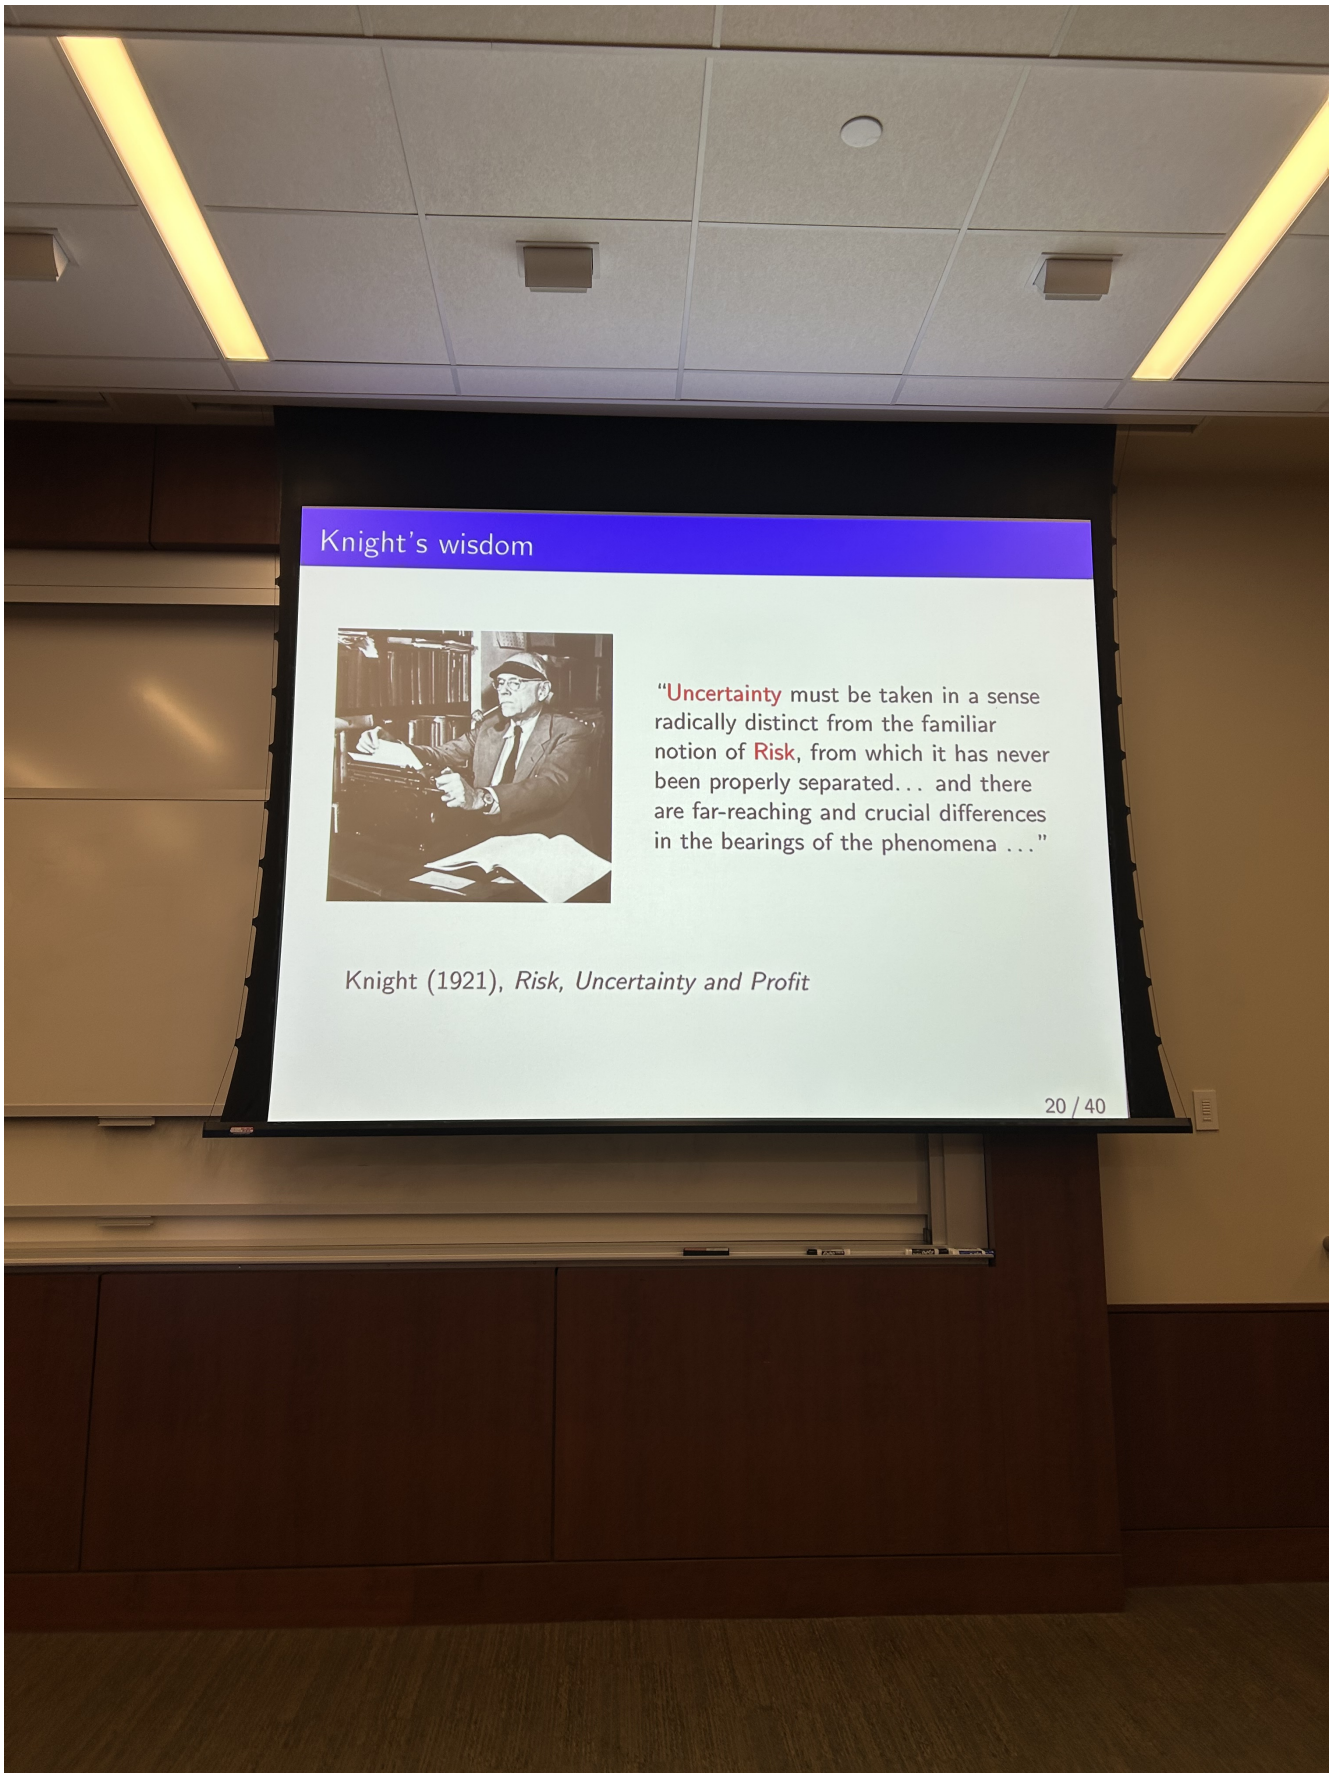


Figure 5: Uncertainty vs risk– a distinction we don't usually discuss in the stat-ML community.

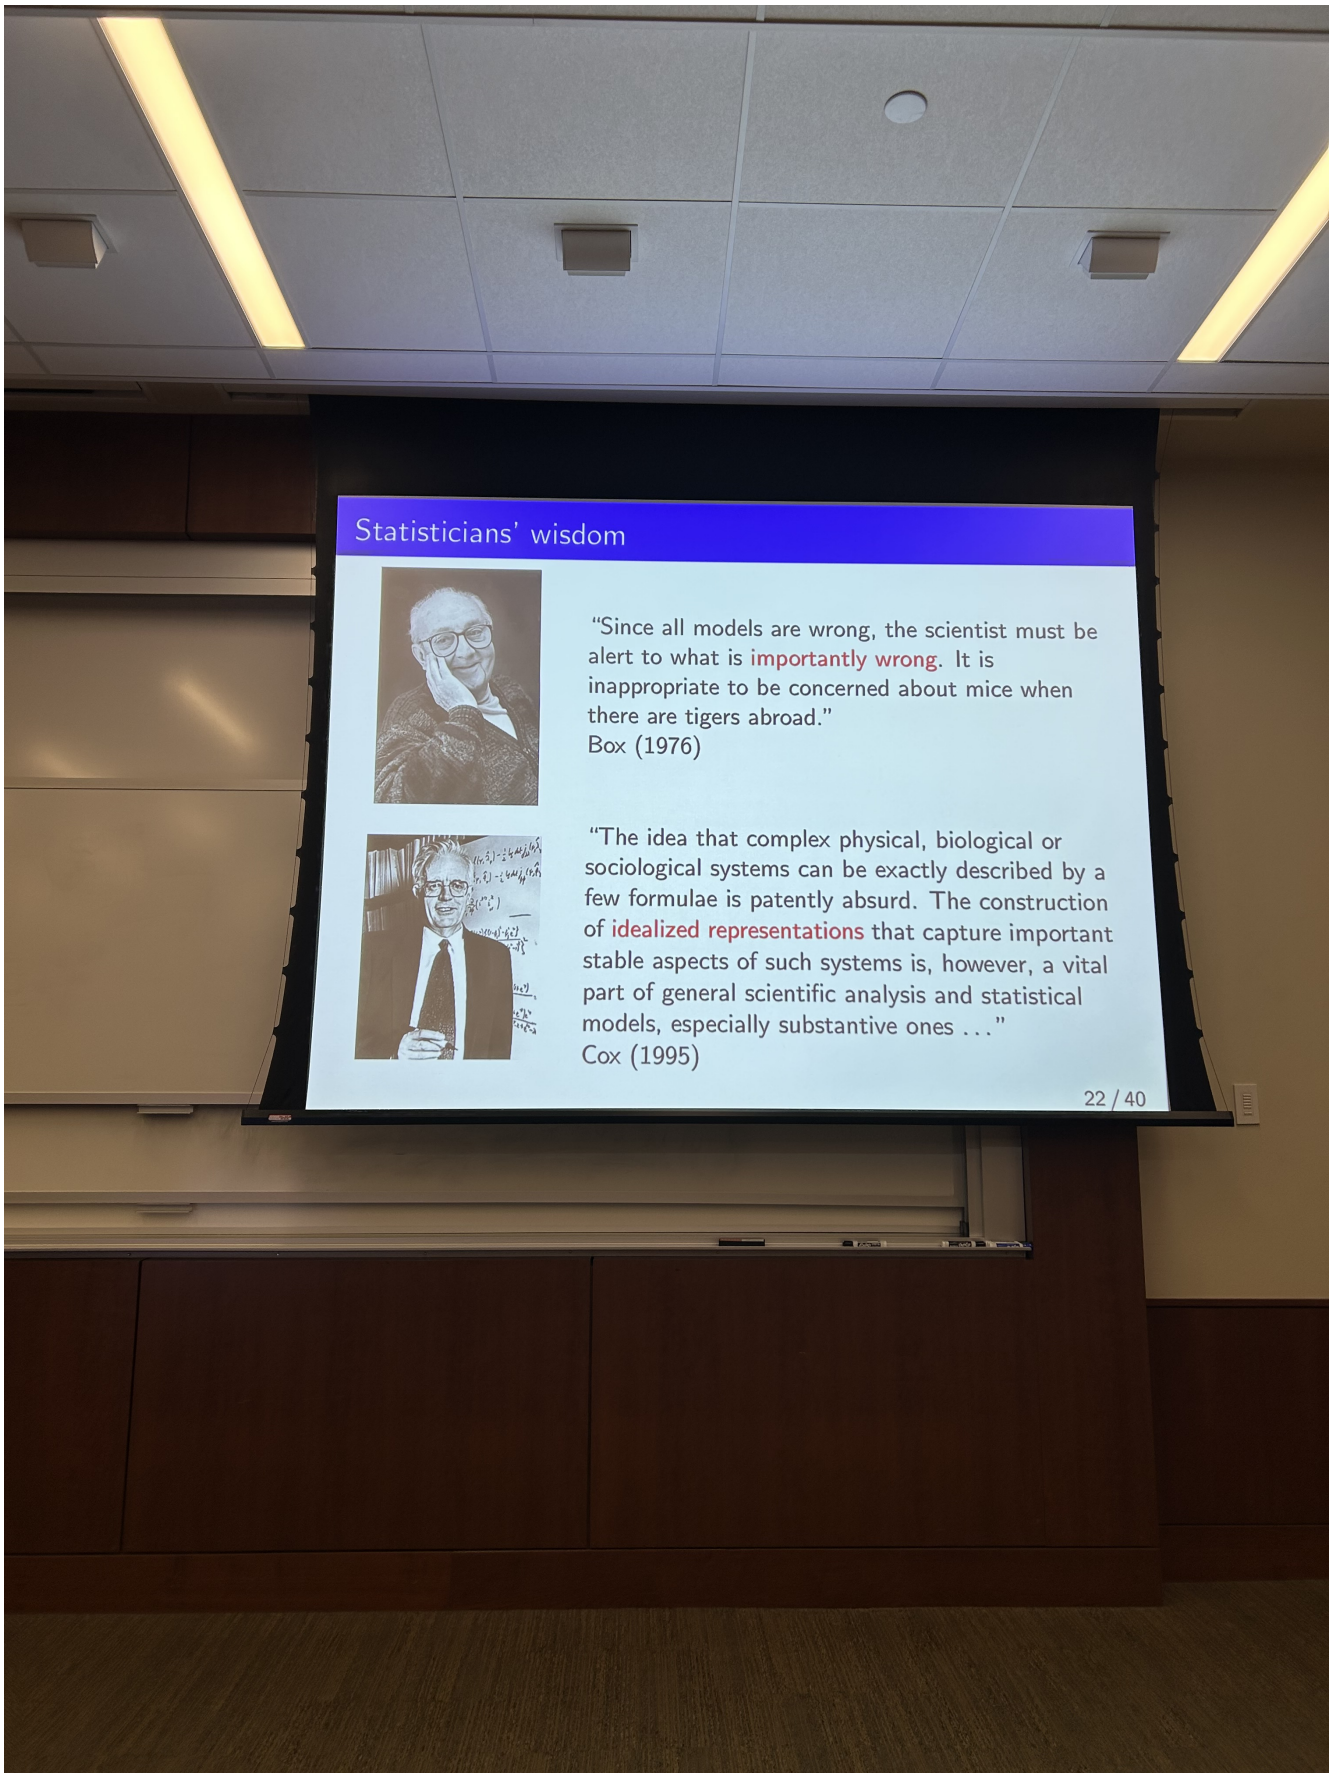


Figure 6: Enter Caption

Outcomes but probabilities are known: risk

Which among multiple probability models is best: ambiguity.

Ways in which a model might give flawed probability predictions: misspecification.

[SJ: Make sure I understand the distinctions.]

"We endow our policymaker with preferences that encode these uncertainty aversions."

[SJ: My guess: it basically becomes a max min robust optimal control problem, as he hinted about above by referring to investment and emission decisions as maximizing controls, and the probability model as a minimizing variable? To confirm.]

[SJ: Q from audience: smooth ambiguity, not Gilboa–Schmeidler (?).]

Basically a sensitivity analysis [SJ: ?].

Misspecification aversion: divergence wrt a baseline model. [SJ: Very arbitrary.]

Ambiguity aversion: divergence wrt subjective priors/posteriors

We invoke the min-max theorem, we investigate the worst-case prior. [SJ: Analogous to the least-favorable prior from minimax theory in statistics.] This distribution identifies the "most concerning" "types" of misspecification. We then use the worst-case adjusted probability to compute valuations.

[SJ: A very skeptical Q from the audience: what is the input? Why is the input reasonable? Reply is clear.]

[SJ: Had to step outside at slide 27/40, came back at slide 29/40; ask Prof Hansen for slides 27 and 28 after the talk; is it just the robust min max problem?]

Part III:

AK model [SJ: ?] for climate change-related "big project" R&D investment and "technology discovery", and discuss investment implications.

Ex: nuclear fusion investment.

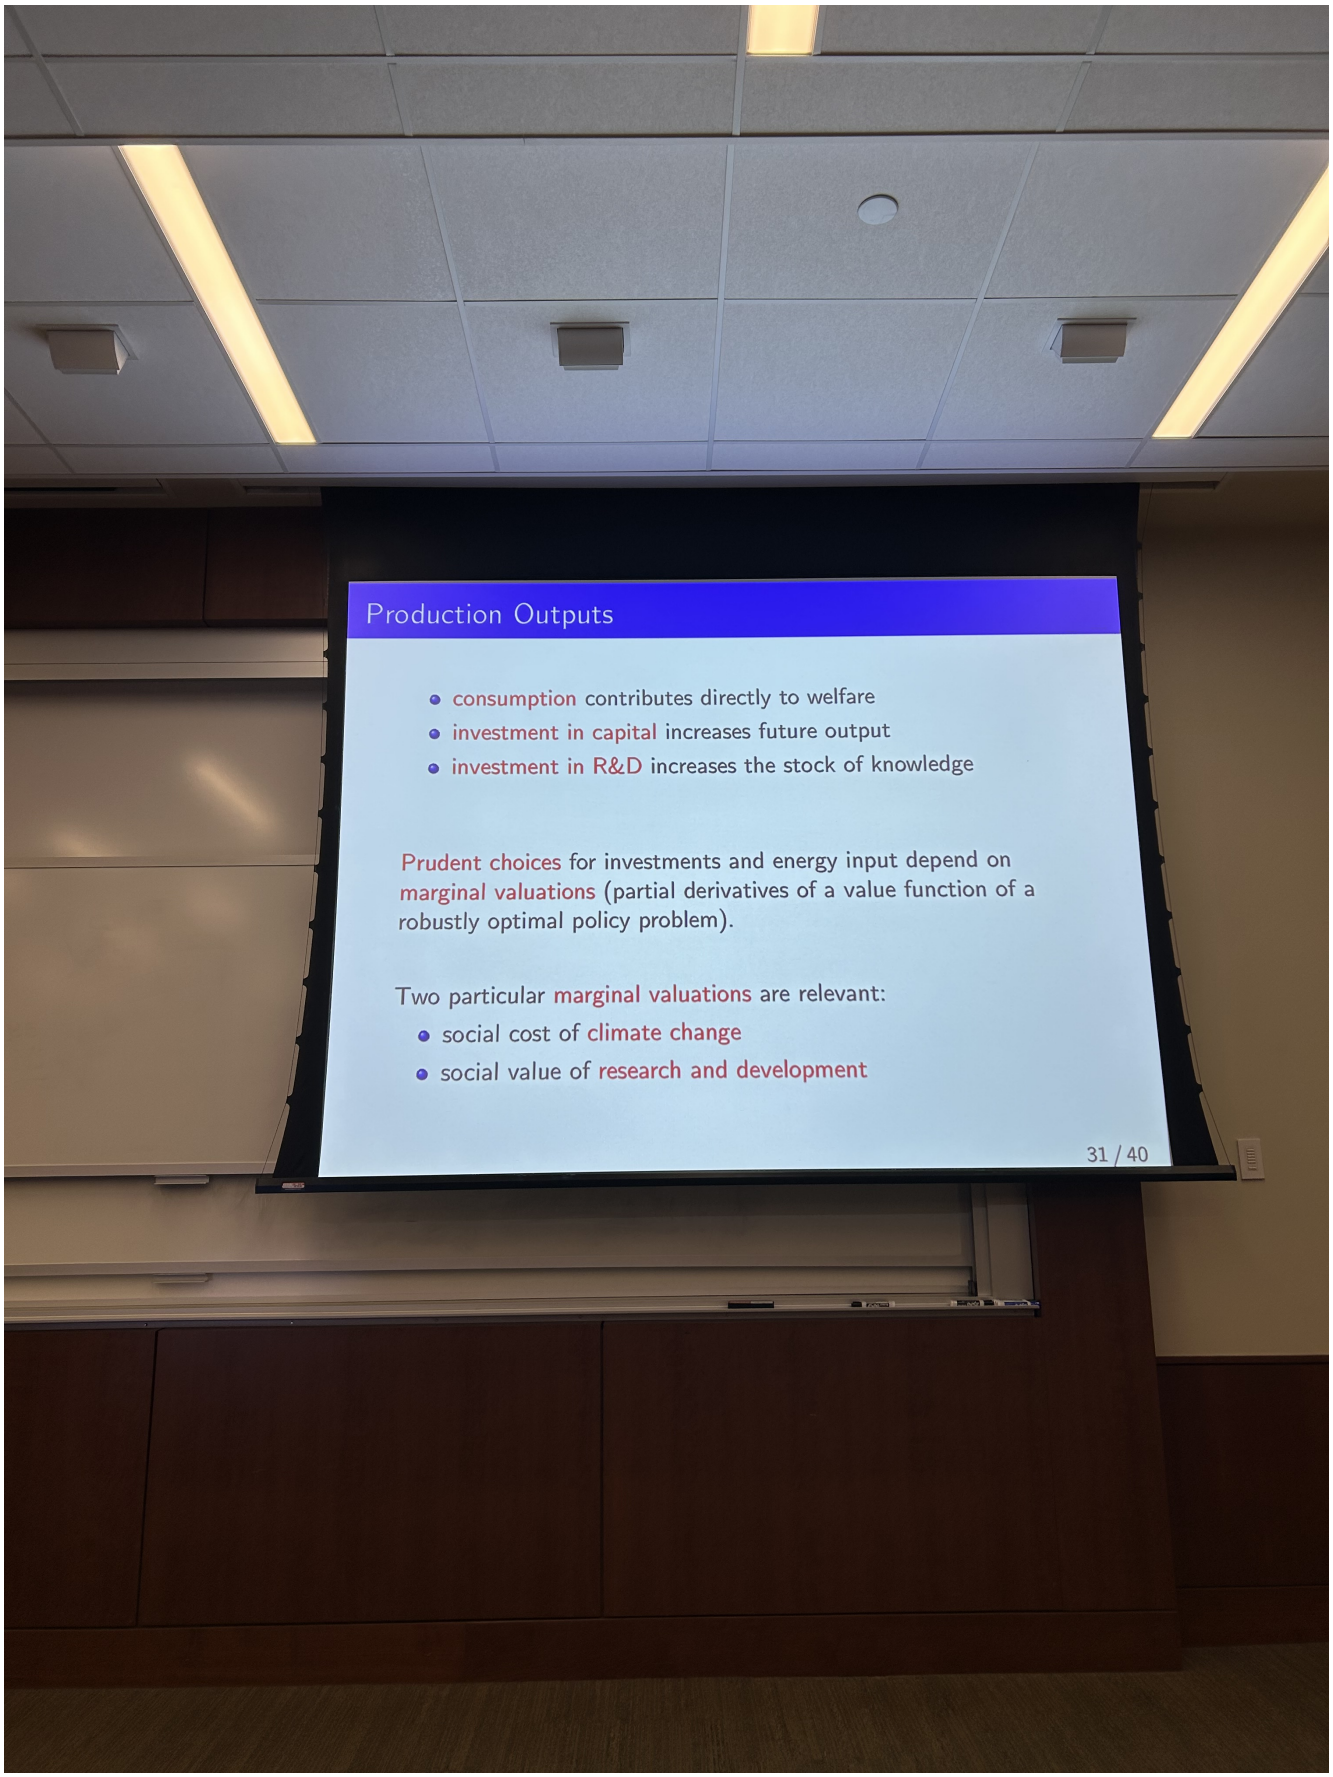


Figure 7: Enter Caption

The technology uncertainty, with respect to economically viable technologies, dominates the uncertainty concern in this model.

See plot of uncertainty-adjusted probabilities for the technology jump only specification; as you crank up aversion wrt the number of years till a technological breakthrough (20, 40, 60, ...), you see that the worst-case prior shifts. (In the model, "arrival of (the first) technological jump" is modeled as the first jump for a particular type ℓ.)

Plot of investment as a function of the inverse ξ_m^{-1} of the "robustness penalty parameter" [SJ: appears in the formula in the min-max problem I missed?]. It's

That's non-intuitive; why isn't it monotone?

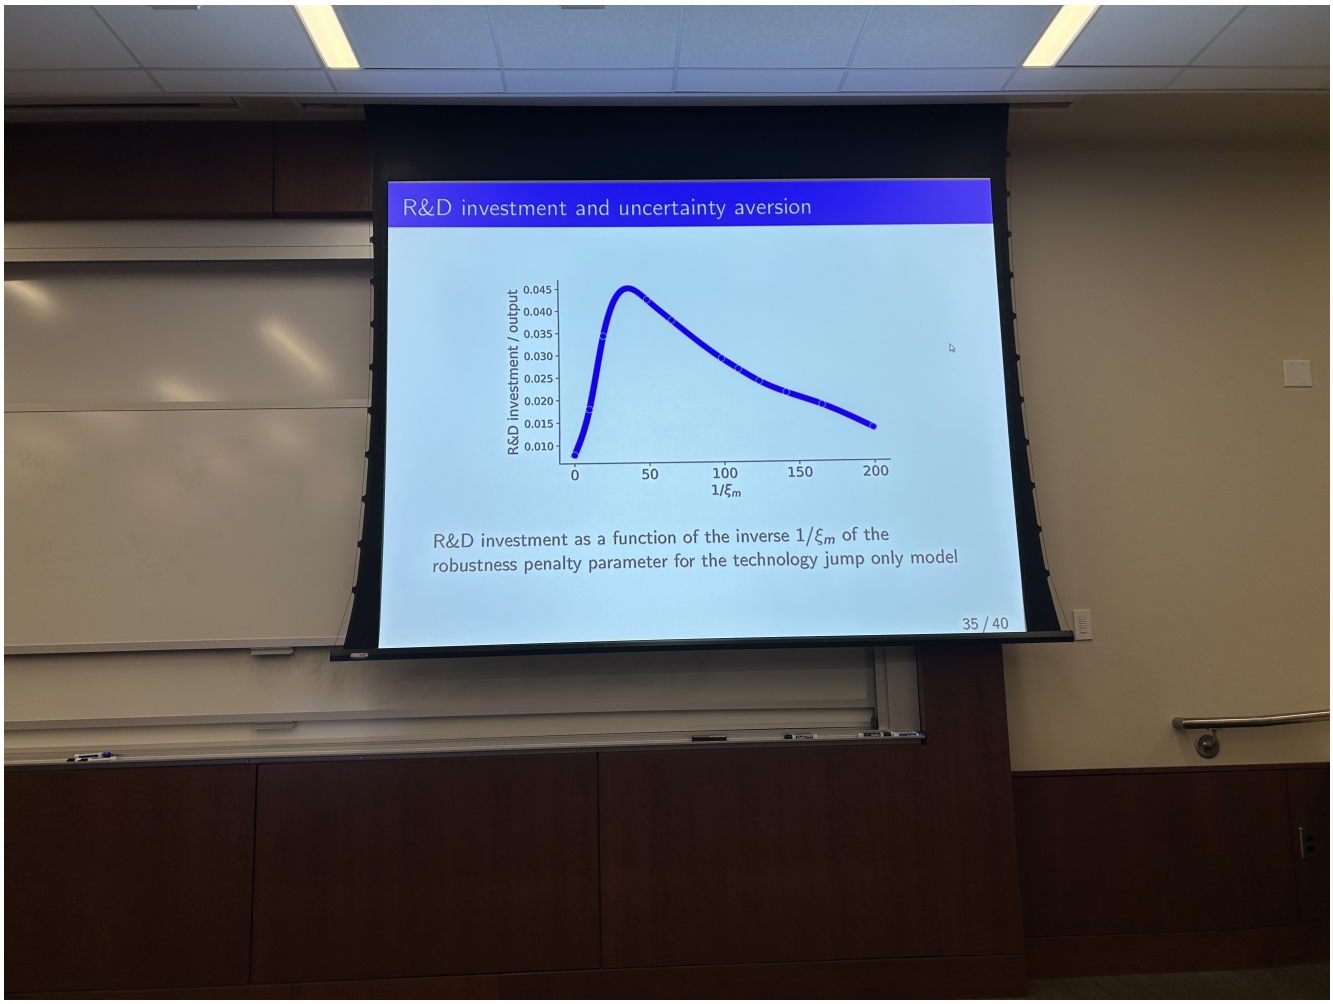


Figure 8: Enter Caption

The punchline (!) is that there's a tradeoff, a fight between two terms in a product, as illustrated in the following slide. [SJ: I haven't digested it yet.]

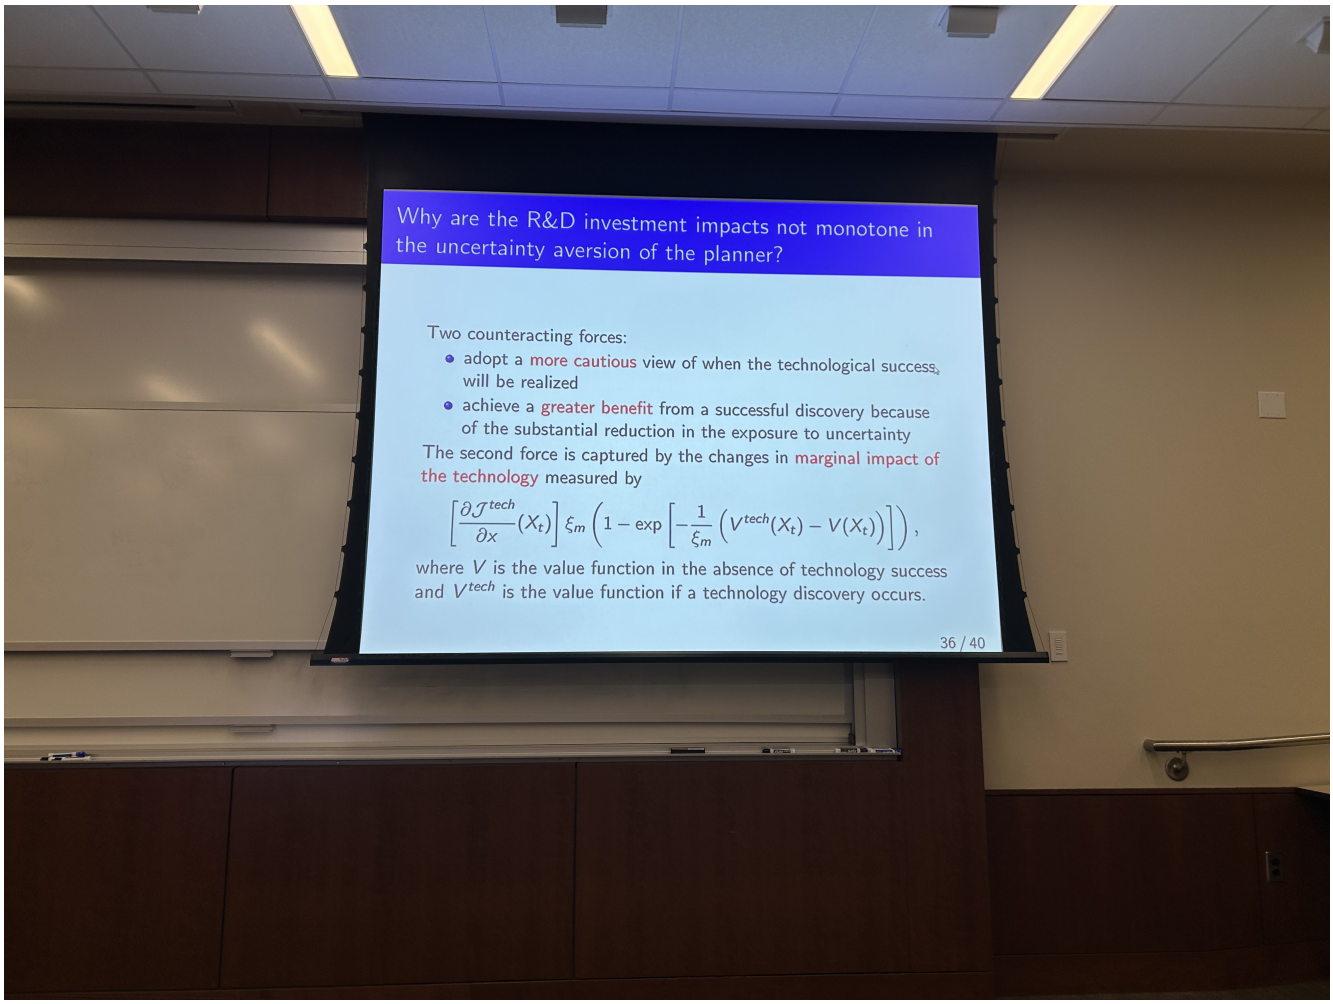


Figure 9: Enter Caption

Follow-up work on mean-field control, mean-field games. Also computational challenges to solving. [SJ: !]

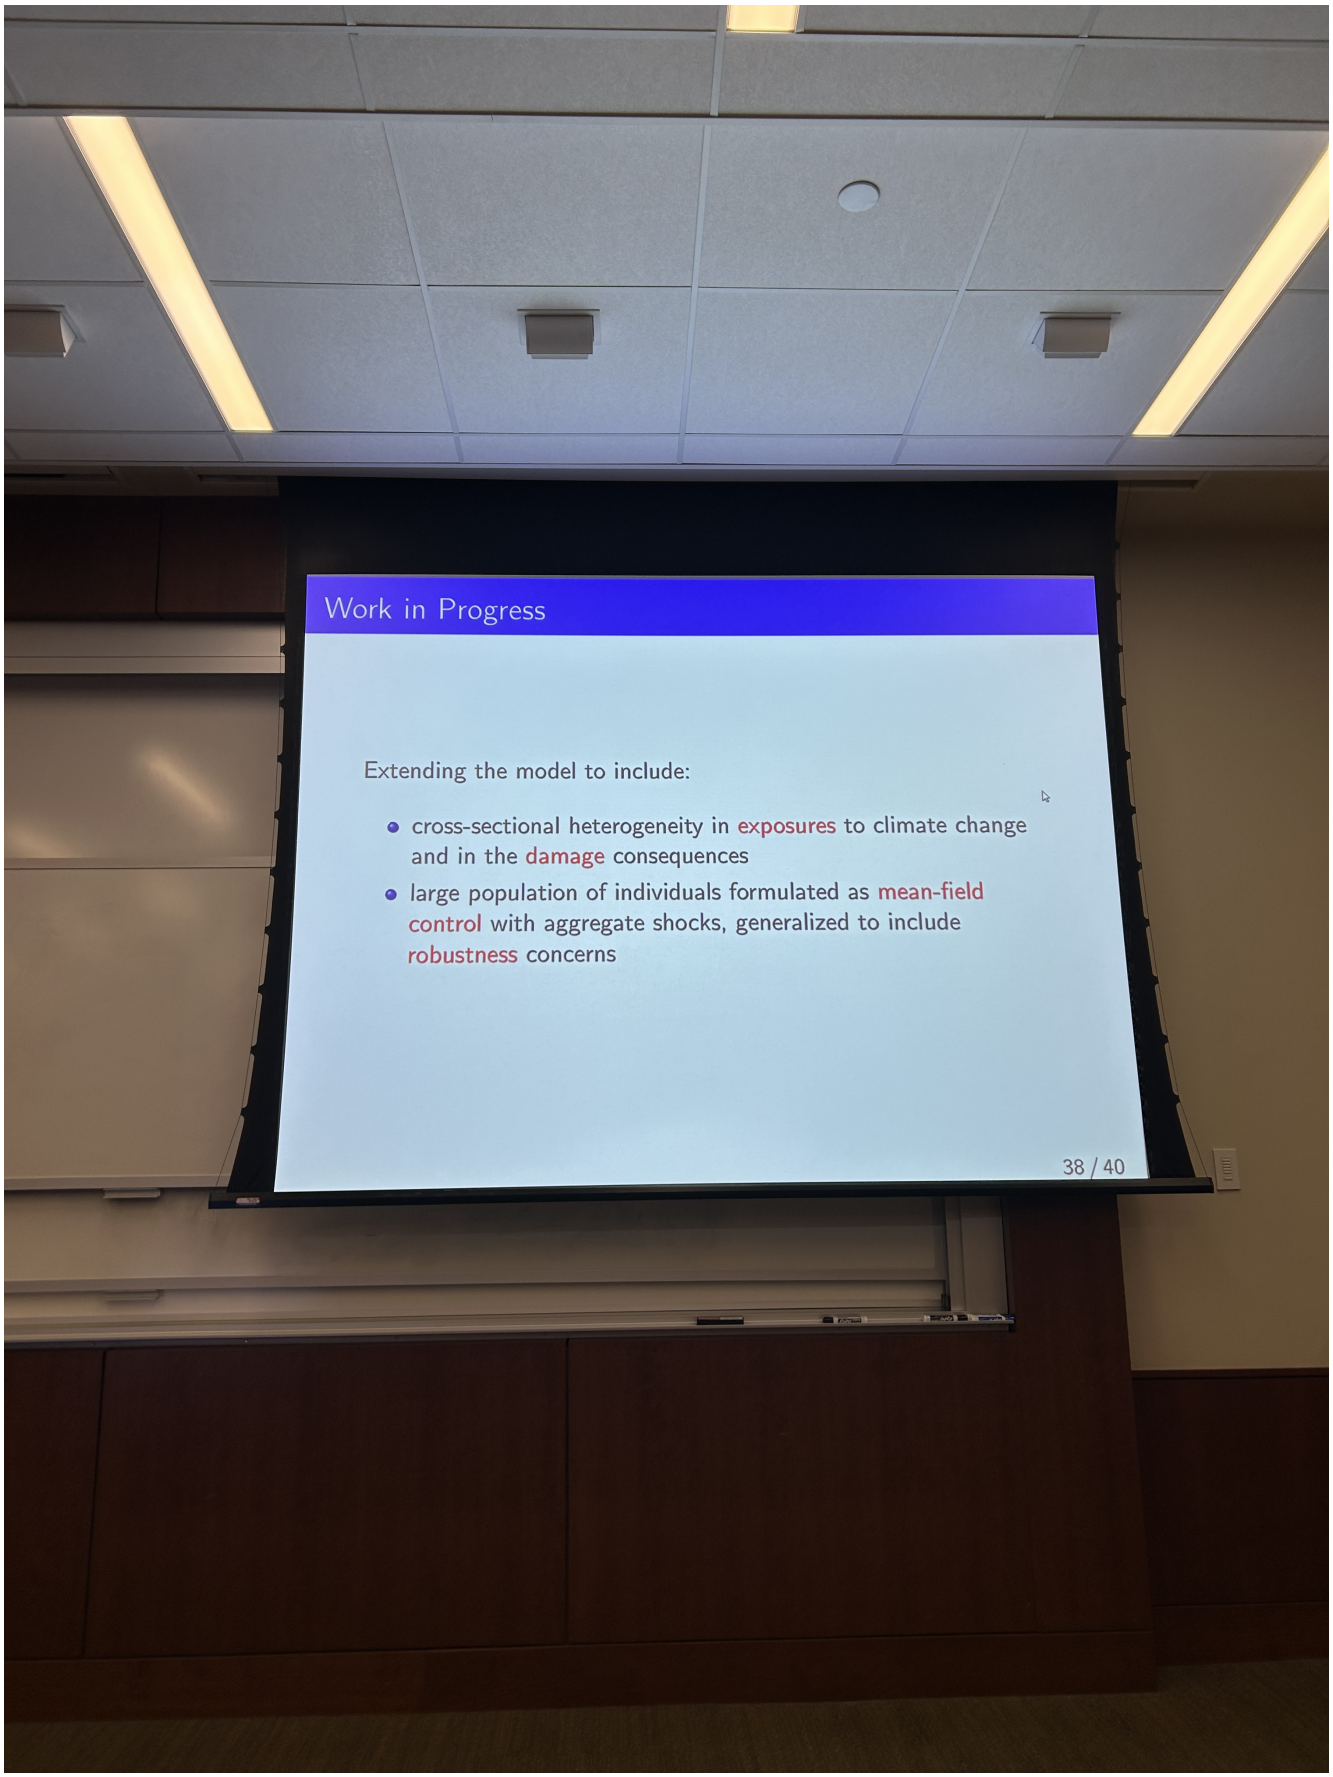


Figure 10: Enter Caption

[SJ: My Q (because I missed slides 27 and 28): what's the formula for worst-case model? Reply: he's doing KL-regularization, hence it ends up being an exponential tilt of the original measure; you could consider other natural divergences, like f-divergences, etc.]

[SJ: My next question: what's the difference between mean-field control and mean-field games? His reply: MFC "couples" the participants in such a way that you just solve a single optimization problem, as opposed to solving a game. It's computationally easier.]