[SQUEAKING]

[RUSTLING]

[CLICKING]

JAMES PAINE: Hi, everybody. My name is James Paine. I'm with the System Dynamics Group here at MIT. I'll give you a little bit more background on myself in a few minutes. But just to be sure everyone's where they're supposed to be, this is System Dynamics-- Systems Thinking and Modeling for a Complex World. This is a three-hour IAP overview session.

Just making sure no one stands up and walks out at this point. OK, good, yes. I'm doing pretty good. A little overview of what we're doing today. So right now, we're in the middle of the welcome, which is nice. We're going to do a quick overview of What is Systems Thinking? And I'm going to try to emphasize this throughout the day, that there is a subtle difference between system dynamics, the method set, and system thinking, the concept and field. And they overlap, and they flow in and out of each other, but they are subtly different.

Additionally, because this is only a three-hour overview, we won't have the time to dive into a lot of the specific tools, and methods, and nitty gritty. Happy to discuss that in more detail, or maybe even I'm thinking about planning a future session with that. So one of the things we thought about-- what's the way that we can get people engrossed in this material as quickly and efficiently as possible? And our thought there is, we're going to do something that we like to do quite a bit which is a management flight simulator. So that's going to take up a good chunk of the middle of the afternoon.

I hope everyone has been receiving my emails. I'm sorry for all of them. But we need essentially one laptop per about three to five people in the entire group. So I think we're good with that. So that will give us an opportunity to get hands-on with the system dynamics-fueled simulator, and then afterwards, we're going to sit down and talk about how the lessons from that simulation tied into some of our earlier concepts. And finally, end with an overview of some additional resources for those of you who are interested not only in system dynamics but specifically System Dynamics here at MIT.

So who am I? I'll say this, I was not realizing that the screen was so big, because my face is very large up there. Yeah. My name is James Paine. I'm a second-year PhD student in the System Dynamics Group. My path here to MIT was torturous, to say the least. I worked for approximately 10 years in various industry settings as a chemical engineer at GE-Hitachi nuclear energy, as an industrial engineer and Lean Six Sigma specialist at a company called Inmar, which does reverse logistics, and then, finally, as a product manager at HanesBrands on the Maidenform and Playtex 18 Hour brands. So all over the place.

One of the reasons, actually, why I found the System Dynamics Group specifically was when I decided to go back into academia as a career. The System Dynamics Group was one of the few places when, during my interviews with different folks, when I said I've worked on everything from nuclear reactors to bras, their response was, OK, that makes sense. And then, my response to that was, well, that doesn't really make sense to me, I'm curious about why that makes sense to you. And then, we kind of dove into it from there. And we'll get more into that sole idea of systems level thinking.

My specific research interests are focused primarily on this concept of behavioral operations management. So some of you in here might be familiar with operations management as a concept behavioral OM that specifically looks at the decision making processes that come about as a result of human beings existing within a supply chain and either trying to work around or with that. My constant line and refrain for all this is, within behavioral OM, the presence of people is not a bug, it's a feature, and what can you do to kind of plan not only around but with that.

So a little bit more about the System Dynamics Group specifically at MIT here. Here are our faculty members. We have John Sterman, Nelson Repenning, Hazhir Rahmandad, and David Keith. John Sterman is a big presence in the system dynamics field at this point. He is the author of-- and I'll talk about this later on-- this lovely textbook right here, which is a big, old sucker-- *Business Dynamics.* This has become the de facto textbook of the field at this point.

There's a number of different resources out there. I'll talk about, again, this at the very end of the class, but given limited time and limited resources, my suggestion would to be grab one of the copies out of the Dewey Library over here and read the first two chapters. That greatly overlaps what we're going to talk about today, and in my opinion, covers a lot of the fundamental concepts of system dynamics and systems thinking. So very brief history of system dynamics.

System dynamics is, at its core, an MIT, sort of, founded field. It came about out of the mind of this man right here, Jay Forrester, and its origins are in control theory. So this is-- I want to say-- I want to preface this by saying this is my view of system dynamics, and this is coming from the standpoint of someone whose background is all sort of engineering and operations management. I've used system dynamics as control theory applied to social systems.

And again, you might have some slight disagreement from different people in the group, but that's at least my one-sentence summary of all that system dynamics is. So the real, seminal work of all this was in 1958 when he published *Industrial Dynamics,* which was sort of the first formalization of these concepts that we're going to be talking about today. Jay Forrester himself-- if anyone here has a background in electrical engineering, you might recognize his name. He's the original inventor of RAM. So his background is primarily around computers.

He was involved in the WHIRLWIND I, which was MIT's first multipurpose, digital computer. So he did a lot of work in digital computing, a lot of work in the electronic space. Actually, if you go up to the fourth floor-- room 450-- it's the Jay Forrester Conference Room. In there, there are pictures of him on his ranch that he grew up on. There's pictures. There's chunks of memorabilia from his history of system dynamics. There's also an original piece of hand-woven RAM shoved up there. So if you have any interest in early computing history, that's kind of a fun bookshelf to wander through and touch and feel.

So, really, how this started was, he came to Sloan and had this idea of what can I do to take my background in engineering and apply it to business systems. And this was born out of his view that, often, when faced with a problem in this space, you could get so far with treating it as a purely engineering problem. And ultimately, you would end up running into social issues, you run into bureaucratic issues, you run into structural issues inside the organization that are ultimately trying to implement the solution that have to become part of the solution in and of itself.

So he asked himself, how can I broaden and apply some of this theory that I've already learned and developed over the course of my life to this wider concept of social systems? His real test of that came in two forms. One was an early project with General Electric at an appliance plant, and they noticed that they had this three-year cycle, specifically with regards to their employment history. It was sort of boom and bust.

Have a whole bunch of people coming on board-- onboarding, onboarding, onboarding-- and then mass quits every three years. And the question was, where is this coming from, what's happening? So he sat down and hand coded one of the first-- what we consider to be-- system dynamics models within this field and presented it. And that led to his work in *Industrial Dynamics-- A Major Breakthrough for Decision Makers.* Where this got a lot of attention and spread beyond our group here at MIT was when he was invited to the Club of Rome and started applying some of this work to larger social systems in his WORLD2 simulation.

This led to a book later on-- I'm sure you guys have heard-- called*The Limits To Growth,* which is, in and of itself, a fundamental book within system dynamics but also a little bit controversial in the form of that it was semi-misconstrued as making concrete predictions for the future. Instead, it was trying to make some comments about the overarching social structures we live in and the limits of resource renewal and regeneration. So that's kind of my quick background of where we are and how we are. What's interesting is, that what that's led to is, a group here that has a wide, wide background.

So these are the current PhD students in our group, with me down there in the middle again. We have two other members in the room. We have Jose back there in the corner. We have TY sitting over here. Just going to call him out. Yeah. But if you look up here, the reason I have this up here partially so you can hunt us all down later on if you have questions. And then, also, for this little line below, where we all came from and how we all got here. And you'll notice that it's kind of all over the place.

We have some engineering, we have some aeronautics, we have electrical. We also have biology. We have geophysics. We have applied mathematics. So folks come to system dynamics through a variety of different means and methods, but ultimately, what-- at least my opinion on this is-- it's from someone taking a systematic approach and noticing that a systematic approach by its very nature must exist within a system and how can we apply that more widely. So just to give you an idea of where this can kind of lead, I'm going give you two examples of some research that I'm currently working on. I'll get back to this later on.

This one is one that I would consider to be more traditional system dynamics, in the sense that-- notice this lovely, little diagram, which we'll talk more about later on-- stock and flow diagram compartmental model. This is applying concepts of systems thinking that we'll talk about a little bit more later on, this idea of nonprofit organizations and decision making within that space, and specifically how nonprofits within the performance landscape in which they operate are subject to unique and different pressures that for-profits are not subject to.

And this is-- none of you should be surprised by that sentence. That's not a crazy thing, but by its very nature, you create this-- and by outcome of this paper-- this performance landscape that says that not only is it more difficult to be a nonprofit in certain settings, it is structurally impossible for you to achieve certain types of goals.

So that's something that's in the whole line with this is like why is it so difficult to be a nonprofit manager under certain circumstances. That's kind of where it came from. Totally different chain of research here is this idea of multi-echelon supply chains and inventory oscillation. This is built on a classic model within this space.

Has anyone here ever played the Beer Game? Beer Game? We got a couple of hands. So this is ultimately looking at the Beer Game and me being frustrated and saying, we've been talking about the Beer Game for 60 years, like, why are we still talking about it? What's going on in this space? So this is taking, again, the idea of people being people, letting their order behavior be what it is, and talk about what can we do for an algorithmic approach to control mechanisms that exist within a supply chain with delays.

So in this case, I'm not really doing a lot with what one would consider to be traditional system dynamics tools, but I'm spending a lot of time in TensorFlow, I'm spending a lot of time in Python, I'm spending a lot of time doing optimization within these spaces. And that still is within this larger concept of system dynamics and systems thinking.

So let's dive in a little bit more. I've been saying system dynamics and systems thinking a lot. And you guys are in a three-hour overview session, so I'm assuming that most of you guys don't necessarily have a strong background on this. Actually, that's a good question to ask right now. Has anyone here taken any of the system dynamics classes here at MIT already? OK, good. Because otherwise, this will be really repetitive for you, so that's great, OK.

So who here has ever seen a diagram that looks like this? Does this look familiar? You have a problem, you identify the problem, you gather data, you evaluate alternatives, you select a solution, and you implement it. Anyone here ever worked in project management? Anyone here ever worked in that sort of environment? Yeah. So my background is Lean Six Sigma-- once upon a time-- and this is a highly-simplified version of the idea of like DMAIC. Though, for those of you who have a background in Lean Six Sigma and know DMAIC, don't jump ahead of me because I'm going to get to something important here.

So within this concept of identifying a problem, finding a solution, and implementing it, ultimately what we're talking about is that we have some sort of goal, we want to fix a problem. We're going to make some sort of decision. By making that decision, we're going to change the state of our system. Great, we did it. Congratulations. I'm assuming-- partially given away by the title of this section of the slides-- you can kind of tell that this is not complete. Also, just by the nature of how I've arranged these variables on this board, I'm kind of giving away where I'm going with some of this.

Because at the end of the day, this is not true in and of itself. This is all embedded in a larger system that matters in the short term. So when you change the state of your system, your decisions have to change. Suddenly, when you've changed how you might implement or control a specific structural process within, I don't know, your supply chain, the way that you go about managing that supply chain on a daily basis is now different. Additionally, because of that change, maybe now your goals are a little bit different.

Before, you were trying to-- again, I'm going to talk in the language of supply chain-- you were shooting for that semi-arbitrary number of 3.4 defects per million opportunities-- which I have a whole other soapbox about why that number is not the appropriate number to choose. And now you've done it, so now you want to do something else, so now your goals have changed. But now, all of a sudden, by you doing this, you suddenly realize that you're affecting not only your one process but other processes, other things within your system. And that can lead to this thing we call side effects.

But however, maybe by me changing the state of my system, suddenly a supplier has to change how they're acting for me. Suddenly, maybe an order stream has to change. Somebody outside of myself and my team has to change their decisions based upon the state of the system. And additionally, that state of the system feeds in right there.

So you can see where I'm kind of going with this. And additionally, when everything is said and done and you've built this lovely process up, with everyone at their own side effects and building off everything else like that-- here's the fun bit-- none of this happens instantaneously. There are delay after delay after delay after delay.

So you, making a choice today that affects your system and achieves that goal within that open loop thinking model, ultimately can come right back around and affect that original input again. This is the fundamental concept of systems level thinking. No decision you make exists in isolation. No process that you affect exists in isolation. The only difference is the boundaries of your system. In fact, I go so far as to say there are no such thing as side effects, there are simply effects that you have not thought about yet.

So what is a system? A system is any set of interdependent parts with a common purpose. Now, that common purpose-- I'll be honest-- is a bit of a caveat, because as you can probably tell from this, if you take the systems level concept and apply it ad nauseum, suddenly you have a system of every atom of the universe, and that ultimately is not feasible. You do eventually have to draw a boundary around your system. So for our purpose, very generally, we're going to say that define a purpose, and then your system is all those parts that exist within that purpose.

And-- I put my animations out of order there apparently-- social and economic systems are highly complex. They are dynamic, tightly-coupled-- and by that, what I mean is, when one thing happens, you can trace it to another thing occurring somewhere else-- governed by feedback. Ultimately, the decision you make affects something that then affects that same decision process in the future. Non-linear, limited information, and there's ambiguity. And typically, they're more complex than human-made, physical systems.

And so, ultimately, when you take a lovely, little diagram, like this of one thing-- kind of water flowing down, water flowing down, water flowing down-- the part that people often forget is that feedback loop that goes right back up to the top again. So when we talk about systems thinking and system dynamics, it's really a framework to help close the loops. I talked about, before, this idea of open loop thinking. That's our method of taking that process and saying, what are the processes that I make downstream, what are the choices I make downstream, and how do they come back and affect that first starting point at the beginning of the day.

Ultimately, what we try to do within system dynamics, we have a general framework for figuring this sort of process out as best we can. A big part of it is eliciting mental models. And mental models in that sense is saying-- walking up to somebody and saying, I know you're acting rationally, I know the choices that you make are the ones that you think are best for the information set that you have right now, I want to know why.

I want to build a model of your mental process that when I run it and I look at it, it makes complete sense, and then take that model and then pop it in the context of a larger system and see how those decisions start coming back and affecting the outcome.

And then-- that's exactly right-- taking that mental model and expanding it outwards and saying, OK, now you've made choices, how can we bring feedback into account? Simulation. So this is one thing-- for those of you who might have a small inkling of the System Dynamics Group here at MIT, a huge chunk of what we do is simulation. Because at the end of the day, it's a heck of a lot cheaper, quicker to simulate a social process, and changes, and tweaks, or impacts, or effects than actually going out there and doing it directly on social systems.

It turns out that takes a long time and often makes people grumpy if you do it directly. So simulation is really, really helpful. So simulation helps also improving mental models by presenting folks with people, also as a teaching tool, showing folks self-contained models and simulations of a process that both shows how their mental model fits and makes sense and is perfectly rational, but then also, in a larger context, can affect and have feedback. So this is the big one. The simulation's purpose is not to be right.

I put this up here partially for my own edification and my own learning. I spend way too much time building these models, both within the core system dynamics tools that we'll talk about a little bit later and also within my own side environments of saying, OK, I want to get this simulation to match some reference mode of data as closely as possible.

And the question there is, though, being right is not the same as being useful. So you want to make sure that you're identifying the high-leverage policy choices within your model, the portions of that multi-loop mode of thinking we talked about earlier that actually affect the outcome that you care about, not necessarily have a simulation that directly matches the physical universe around us.

These are next-- what I consider to be some of the three core points of systems thinking and systems dynamics, and this is the big one. Structure generates behavior. If you walk into a system dynamics conference, if you walk into any room with enough people who do system dynamics for more than five minutes, eventually someone's going to say this. Structure generates behavior. And I'm going to spend a little more time on this in a few minutes.

The whole idea of structure generates behavior is that the actions that people take are possible, feasible, and rational because of the universe in which they exist, and it's important to realize that. And that actually then goes into that mental models matter a lot.

So physical structure-- when we talk about structure, it's not just the physical structure of your system, it's the structure that exists within people's heads. It's the choices that they make when they take an input and then do an output. Again, my background is all process control and process engineering. So one way you can think of this as little tiny PID controllers floating around in the universe taking inputs and outputs.

However, figuring out what those inputs and outputs are and how people convert that information is difficult and should not be casually assumed away. And then, finally, the fundamental attribution error. This is something that I added back in here, because this idea that our first instinct is to blame the people and not the system-- and this goes back to this idea of structure generates behavior. So like, for example, if you're on the highway and suddenly someone cuts you off, what's your initial reaction?

Screw that guy! What the heck! Honk, honk, honk. Now suddenly, you're on the highway and-- I don't know-- you're late to get somewhere-- for me, I got a call that my kid is throwing up at preschool, I got to get there in the next 10 minutes-- what do I do? I cut somebody off saying, sorry, it's OK, and I zoom off. So in that situation-- that first situation, when I'm honk, honk, honk, that guy-- that guy who just cut me off-- what's wrong with them? Something is wrong with that person. No. There's an underlying reason, there's an underlying rationality behind what they're doing.

I will argue that there might be rationality or that person is just really bad at driving, but hopefully, that's not it. But the idea that the fundamental attribution error is that our first instinct is to blame the people and not the system. Back to structure generates behavior. It is the structure of the system that matters less so than the behavior. And even within our own group here at MIT, we have to remind ourselves of this idea of the fundamental attribution error. And we try to break away from it, we're not perfect at it.

This is written on the whiteboard up in our room. This is one of the first things I saw when I walked up to that group, and at the time, I just thought it was a nice saying and I didn't realize this was actually addressing one of the fundamental decision-making fallacies that we have to overcome as system dynamisists. So this is the basic assumption. "We believe that everyone in this community is intelligent and capable, cares about doing their best, and acts with integrity, and wants to learn."

So the assumption is that, every person you meet, every person you sit down with has this operating in the background and you have to assume that they are rational, caring human beings who are doing what they think is best. And the moment that you accept that, then, suddenly, it's not the person you have to change, it's the system in which they're operating that needs to get shifted. So this some more, kind of, burning into this whole structure generates behavior.

You have patterns of behavior-- so events are the most visible. Underneath events, you have patterns of behavior. Under that, you have structure. So when you look out into the universe and you see something, those are the events. You see those things happening over and over again, those become patterns of behavior. And then, the system under which it's generating it is the actual structure.

So here's an example. This is all about oil prices. "Drunk trader caused a spike in oil prices." "Oil prices keep falling-- this is why." "OPEC rumors continue to pull oil prices higher." OPEC turns for high oil prices. "Another sign of economic worry-- tumbling oil prices." Oil prices after tanker attack in the Gulf of Oman. So in this case, each one of these ones, you notice they're talking about oil prices jumping up, oil prices going down, oil prices jumping up, oil prices going down. These are the events that you see. These are the points of data. The pattern is a little bit more interesting.

You take each one of those effects, you can't just draw a line out ad infinitum, you want to look at something like this. So in this case, you can see the oil prices-- and I would see this as sort of a stereotypical boom and bust pattern. It goes up, it goes down for quite some time. Goes up, and then goes down. Goes up, and then goes down. I would actually even boil in here and say that it looks like the amount of noise around this up-and-down pattern is increasing over time, which in and of itself is a little bit interesting.

But ultimately, it's that structure in the background that matters. There's the physical structure, there's the information availability, and then there's also the actual mental model of the actors involved in the process. So we talk about this. This is a whole thing. Whenever you talk about mental models within this space, one of the things that's easy to say is, OK, well, you tell people about this. You tell them to consider feedback, you tell them to incorporate this information around them, that that'll fix it, that'll fix it.

No. It is incredibly difficult to learn and change in a dynamically-complex environment, and these are just some of the information-- some of the reasons why. I'm not going to walk through all of these. I think everyone in this room, probably, is on the same page with me on this one. One thing I'll point out, right there, is limited information and time delays. Gigantic reason. People can only make the decisions given the information they have available to them in the structure in which they're able to operate.

To expect someone to operate otherwise is to expect them to be omniscient in some way, shape, or form, and ultimately, that is not feasible within a larger system. So system dynamics-- in my mind, the System Dynamics Group here at MIT is really applied systems thinking. It is not a method. It's not a model necessarily. It is taking the concepts of systems thinking and applying them in ways where we can get research outputs that ultimately improve some sort of social system for the world.

So now having said all that about systems thinking, I'm now going to dive into some specific, actual tools that we can go ahead and use here within this space. Let me make sure how we're doing here with time. We're doing great. All right. So systems thinking, ultimately, and system modeling is iterative approach. There is no right answer at the very beginning, there is only the answer that you think is good enough to solve or address the problem you care about. It's a spiral approach.

Just like everything else we'll talk about with open loop thinking, the process of building a system dynamics model is an attempt to close the loop. And it's not just compartmental models. So when people talk about system dynamics models-- if you hear that phrase-- what they're likely talking about is this lovely thing over here. This is from one of my papers, and I'm working on it right now. We'll talk about this in more detail in a second. It's called a stock and flow diagram or compartmental model.

This is often confounded with system dynamics, the field. It's partially because this is the go-to modeling choice for many people in system dynamics, and this is also the first modeling choice back in the '50s, '60s, and '70s when approaching these sort of systems. And then, of course, that says all models are wrong, but some models are useful at the end of the day. So when you're taking a spiral approach to systems modeling, what you want to do is you pick something within this setup. You think, OK, what's a reference mode.

So by reference mode, I mean what is a mode of behavior that you can observe and pin down. This person made this decision when exposed to these inputs and had this output. Great, we have a reference mode. Can I build a model that reflects that reference mode? Awesome. Now that I've done that, let's go ahead and add a new reference mode. Can I make this reference mode exist within my existing diagram?

Nope. Let's go ahead and add more structure. And as you add that more structure, maybe it comes back and it feeds back into itself and you have to go back and kind of rework from the very beginning. This is where this concept of simulation becomes really, really helpful, because when you interview somebody-- as I've done with some of the work in the, especially, non-profit space-- you might get six different stories describing three different outcomes across two different possible modes of behavior.

And the question there is, how the heck do I incorporate this all into a dynamic model at the very beginning? My argument there is, you don't. Pick one. Pick one to model it, then when you're done, pick another one and model that one and then see, OK, how do they overlap, how do they fit into each other, how can you make this one, unified model?

And the big thing is, as you're going through this whole process, doing your sensitivity analysis and you're testing, you're going back through and updating your own assumptions and mental models about what you think the problem was to begin with. This idea of spiral approach.

So, OK, now we're actually going to jump into some actual-- I said before that system dynamics are not compartmental models. I'm now going to teach you guys about compartmental models because we use them a lot in system dynamics. Heh. So this is fun. So causal links. A causal link is a fundamental tool. So causal link is simply saying, if one thing changes, another thing changes. This is super simple. So let's start with this one.

You have production, you have inventory, and you have shipments. If I increase my production of a good, what should happen to my inventory? This is not a trick question.

AUDIENCE: Goes up.

JAMES PAINE: It's going to go up. Woo! If I increase my shipments of a good, what's going to happen to my inventory? It's going to go down. You guys are-- woo! You guys are great. So there you go. So these are causal loops, or causal links, similarly. Let's say I have orders booked-- and now, you could argue that there's some missing components in here, but we're keeping this really high-ended.

If I have more salesforce in general, what do you think would happen to the number of orders I have booked? Increase, yep. If I increase my price, how many orders do you think I'm going to get? They're going to go down, exactly, yeah. Pretty straightforward. One more and then I'll leave you guys alone with this one. Population. More people are born, it's going to go up. More people die, it's going to go down.

Great. Causal links, straightforward. Where it gets a little tricky is when you have an ambiguity. If all of a sudden you say to yourself, well, you know, salesforce, I don't know about that. Maybe if I hire too many people with my salesforce, then suddenly they start conflicting with each other. Maybe they start going after each other. Maybe they start overlapping, and maybe customers get annoyed because they get too many calls, so maybe it will go down.

The moment you have this idea that that causal sign could be a plus or minus, that implies that there's something in the middle. That means that that loop does not exist in and of itself, there is structure that you have missing that you need to add. Now, something also, too, kind of the same idea-- causal thinking. And some of you guys might have seen this already. Causal thinking is truly causal thinking. It needs to be causal, not purely corollary.

So ice cream-- so this is generally true-- as ice cream sales goes up, so does the murder rate. So next time they open up a new Baskin-Robbins, we should maybe change towns. I guess that's what that means. Now, as you guys can probably guess, what that really means is that there's something missing. This is not truly causal, instead we have the average temperature. As the average temperature goes up, people start interacting with each other more, they start going outside more. Just the amount of incidents of people kind of interacting goes up, murder rate happens to go up. Additionally, people want more ice cream.

So there. In that case, we have something in between. So something to keep in mind when you build a causal loop diagram is that word causal, when you're putting this together. And then, if you have time, I would encourage you guys to go just like, I'm going to try-- I'm going to run these slides by the faculty and make sure there's nothing in here that I'm not supposed to share, but assuming I can, I'm going to send this out to everybody as well so you'll have access to these slides.

One of my favorite websites, Spurious Correlations. They have all sorts of fun little things like that. So, now, bringing it all together into a loop. So in this case, we have a basic loop of employee skill versus customer satisfaction, complaints, manager's time spent resolving customer issues, and manager time spent coaching employees. I'm going to walk through this really quickly. Again, we only have a few hours to talk. So normally, this would be a little-- if you guys end up taking one of the system dynamics classes, we step through this very carefully. Yeah.

AUDIENCE: What do you mean by causal?

JAMES PAINE: So causal, in and of itself-- you raise an excellent point. Causal, in and of itself, is the degree to which you have confidence on it. So at the end of the day, causality is the realm of mathematicians and physicists, in my opinion. True, true unresolved causality. So causality in the sense that, if one thing occurs, the other thing if, and only-- one thing occurs if, and only if, the other thing occurs, that they truly are connected in time and space. When one happens, the other one follows.

Corollary is that that may still hold true, but the reason they're happening is because of an in-between mechanism that is driving both. So where this gets kind of tricky is, the lines of causality become blurred when you start talking about macro-social systems. And ultimately, this is the same idea that all models are wrong, but some models are useful. There's a certain point of time where you have to pump the brakes and say, all right, this is causal enough for my purposes and I'm going to move on.

But the same thing-- when you said the moment you have a time where that loop polarity can shift or, suddenly, something increasing causes another thing to increase. Now, suddenly, that same thing increasing is causing everything to decrease. Those are no longer causally-related in space and time. There's something in between. There's some mechanism in between that's causing that thing to go up in this direction causes this other thing to go up and/or down depending upon some in-between circumstance.

The moment that degree of specificity is necessary, then suddenly that causal link needs to get split out into more detailed structure. Maybe I'm diving into too much detail. Is that helpful? Kind of?

AUDIENCE: [INAUDIBLE] you mentioned physics, but like there's a perspective of physics in which you actually look at the physics equations, there is no real such thing as causality, there are just states of the universe that are mathematically consistent and states of the universe that are not, so causality-- so it's like, causality really depends on the assumptions of your system. Like, you're assuming that things could have been different, but in reality, they couldn't have. So I guess I was-- the operational definition of [INAUDIBLE].

JAMES PAINE: You're completely right. Is your background either physics or mathematics?

AUDIENCE: Yeah, physics.

JAMES PAINE: Aha, there we go. So yeah. No, no, you raise an excellent point, because I had the same conversation with a couple of other folks in the physics group. And within our-- what it boils down to-- within our social system modeling, our definition of causality is loose at best-- but ultimately, it boils down to, if one thing moves in one direction, another thing moves in the same direction consistently every time. The moment that you can't say the phrase consistently every time, then that concept of causality breaks down. Just so you... Yeah.

AUDIENCE: So when you're dealing with social systems, does it become a little bit less [INAUDIBLE] and more statistical?

JAMES PAINE: Yeah. To be very honest with you, I have to be careful because I think some folks in the group might disagree with me on this. But again, my background coming from control theory and operations management is that the realm of causality and the realm of correlation are degrees of convenience at the end of the day. This is a gigantic asterisk in the top corner. This is the opinion of James Paine. Please do not necessarily tell other faculty that that's me.

But yes, it is a division of convenience. And ultimately, what it boils down to me, is that, if you shove something in one direction, does another thing constantly move the same way every single time? The moment that it doesn't, then they're no longer coupled within these systems. That's the very, sort of, hand-wavy, simplistic way of approaching it.

AUDIENCE: All else being equal.

JAMES PAINE: All else being equal. That's a good point. All else being equal. So here. I'm going to go through these loops real quick. This is good. So this, I think, is pretty straightforward. This is the whole idea that, as your employee skill-- this is, again, very general-- employee skill goes up, your customer satisfaction goes up. Customer satisfaction goes up, complaints go down. Times that your manager now spends resolving customer issues goes down, so now the manager has more time to coach employees.

So guess what? Employee skill goes up. This is an example of a causal loop. You start at one variable, yet work your way all the way around to the same variable again. Now, I'm speeding up here just a little bit for the sake of time, but let's just go ahead and say that we somehow exogenously increased the amount of time that managers can spend coaching their employees. I don't know, we somehow give them extra time in their day. What does that mean?

Well, this would imply that our employee skill you'd expect to go up. And if our employee skill expects to go up, then we expect our customer satisfaction to go up. We'd expect our complaints to correspondingly go down, based upon these causal diagrams. If our complaints go down, then the manager time spent resolving issues goes back up, which suddenly means that the amount of time they have available to them goes back up again.

We've gone all the way back up the loop, we've increased one variable by a little bit, and we've gotten a positive gain around that loop. If any of you have any experience with control theory, this will look familiar. This is a positive gain or reinforcing loop. This is a loop that, essentially, runs away in one direction or another. When one variable starts going a little bit up or a little bit down, by its action, it causes a cascade of events that causes it to go more up or more down within that space.

Correspondingly, we can have something-- I think you guys can tell where this is going-- called the balancing loop, to get ahead of myself a little bit. So in this case, the general attractiveness of your market goes up, the number of competitors within that space-- you could imagine-- would go up. Your product price-- this is due to competitive interactions-- might go down. This is an example of an aggregated causal diagram.

I think you can make a really strong argument that those are not necessarily perfectly causally-related, but for our conversation, it's good enough. Product price goes down. Your profits, all else being equal, you'd imagine would go down. That's-- OK. I think I did that one a little backwards. We'll find out. Anyway, balancing loop. So the whole idea with the balancing loop is that a small-- these are loops where they tend to balance themselves out.

You have some sort of external stimuli that pokes a variable up or pokes a variable down, it will act in a way to return back to its initial position. A good example of a balancing loop that comes up a lot is called the goal-seeking loop. In this case, you have some sort of desired performance out in the universe and your system will act to close a gap relative to that. I'm showing you guys a structure kind of relatively quickly to give you some chunks of system dynamics compartmental models that you might see frequently.

So the right way to tell if you have a balancing or reinforcing loop is you trace the behavior, like we did before. You pick a variable, you kick it in a direction-- plus or minus-- and you see what happens when you loop it all the way around. There's a quick way to do it, which is count the negative. If there is an even number, then it is a reinforcing loop, zero or even. If it is an odd number, then it's a balancing loop.

The part where that gets a little wonky is that if you mislabeled it in some way, shape, or form, then suddenly that no longer makes sense. So that's my little asterisk in the top corner there. Stocks and flows. So stocks and flows, ultimately-- I'm going to skip ahead a little bit here-- are anything that-- a stock is anything that has memory within this space.

So if you go back to the origins of system dynamics and its applications in control theory, a stock is anything whose value depends upon the instantaneous or the accumulation of change of two other variables.

So how we tend to show this in our space is as a bucket with some pipes. So you have an inflow coming in, you have some sort of accumulation in the middle, and you have some outflow coming out of the back end right there. This right here, now, allows your system to have memory. This is what your system has when you talk about number of employees, you talk about number of units in your inventory, something else that has sort of stickiness in time, something that instantaneously doesn't disappear or change within one unit of delta time but rather exists and persists.

So we use something called the hydraulic metaphor. This picture-- the quality of it is really low, and I should remake it. It's the idea of a bathtub. So you have your pipe coming in, you have your pipe coming out. And ultimately, you can't get rid of something in the middle unless you take it out of the back end or turn down the amount of inflow on the front end. So a good example of this is greenhouse gases in the atmosphere.

What do you do about it? There's really only two options at the end of the day. You have to either increase net removal or decrease net emissions. Those are your two choices. So that's the one thing we talk about. A stock is something that accumulates over time. So this is a good example right here. This is just some stocks versus flows. Your balance sheet versus your cash flow statement. Wealth versus income and expenditures per unit time. CO_2 emissions, vehicle production-- those are all flows.

Part of this-- you can tell-- is based upon how you define the variables at the end of the day, so you have to be a little bit careful. But I will give you this. Here's one that I kind of consider a little bit tricky-- interest rate. What would you consider to be an interest rate, a stock or a flow?

AUDIENCE: [INAUDIBLE]

JAMES PAINE: I called it tricky, so you can probably think, like, oh, take what you think it is and then just flip the answer. That's what I would do. So it has the word rate in it, which immediately makes me think flow, but it's actually a rate on money-- it is a price. It is something that says that, if you give me x number of dollars, I will give you back x number of dollars over some period of time elapsing. In that case, it is a price, it is a sticky quantity, it is a stock, so the word rate there is misleading.

I'm going very quick guys because I want to give you guys time for the simulation. So that right there is really quick system dynamics in a nutshell, in terms of both systems thinking and concepts and also some of the fundamental tools that we use within stock and flow diagrams. What I wanted to do next was transition into actually a very large group project that will give us some time to really experience some of this firsthand, and then ultimately, tie it back in. But before I do that, I want to give us a few minutes to talk about this.

Anyone have any questions about either System Dynamics Group at MIT-- and we'll have some more time at the end of this-- or system dynamics and systems thinking as methodology at this point? Yeah.

AUDIENCE: When do you think-- or when is a problem too complicated for you to say system dynamics is useful? What are the bounds on usefulness of applying this to a problem?

JAMES PAINE: Well, the weird thing is, I won't necessarily say-- complicated is not the right word. I would say that system dynamics is a poor predictor of the future when it comes to precision. So like, there's an example of system dynamics, I think, is good at making models that predict modes of behavior but not necessarily point predictions in the future. So like, a good example would be an oscillatory mode.

If this is your idea of-- this is like your mode of behavior of something over time-- I don't know-- this is like some price of something over time, and you make a model that ultimately shows that, given some period, some set of inputs, I can do something like affect the amplitude and periodicity of this scale but ultimately does show something that kind of changes with time, but for some reason my model is just a little bit off.

The question is, which is useful-- knowing that I have this sort of up-and-down amplitude and here are the things or features that either predict a change in that amplitude and frequency, or is it more useful to know that the price is going to be this versus this at this time in the future?

And the issue with any sort of systems model is, the moment you move far enough away from whatever your current sort of stab in space is, you can be just as wrong as you can possibly be when it comes to point predictions in the future.

So my argument is that system dynamics is good at coming up with models where you can start looking for these modes of behavior and identifying the chunks of structure that have a large policy impact on affecting or adjusting these modes of behavior, not necessarily figuring out am I going to be here or here at time this. So that's-- yeah. I got to use the whiteboard, so that's fun.

OK, cool. So what we're going to do-- this is another bit of bread and butter of the MIT System Dynamics Group-- a management flight simulator. This is the whole idea of working on mental models. Ultimately, getting someone's hands dirty with something like this and getting an opportunity to experience it firsthand is one of the best ways to teach concepts from system dynamics and systems thinking. So *Fishbanks.*

Fishbanks was originally developed quite a number of years ago as a board game-- and I think you guys will figure out very quickly what it's trying to elicit out of you all. This is a smart group of folks. But it has now been turned into kind of this interactive simulation.

So for the next hour and some change as we go through this-- I know everyone here has different backgrounds, different businesses, different things you've been exposed to, but this is going to be your business. Your business is going to be deep sea and coastal fisher people. Your job is to go out there and catch fish as best you can and make as much money.

You are all endowed with your own small fleet of boats and your own commercial fishing operation. Congratulations, welcome aboard. So that's all I'm going to say about *Fishbanks.* Again, there's a lot more stuff to say about this, but I'm going to go ahead and wrap this up, because, again, we only have three hours in total. Before I do that, any final questions about the simulation itself? Any accusations that I rigged the whole thing and made it impossible? Or anything fun like that? Yeah.

AUDIENCE: One of the major, like, one of the teams also found out that their change happened too fast, and I think that's because we didn't understand the system or the equations behind it, the model. Was that information left out intentionally?

JAMES PAINE: To a degree. So one interesting way to run this simulation-- and it's been done, but it hasn't been done for a long time-- is to essentially run this whole thing again with you. Like, essentially let every person in this room play the game again, but understandably, with slightly different parameters, like so you can't just set it straight to what you want to do right off the bat. So the argument here is that information is imperfect, but it is still present and some information is better than no information.

So the constant question is, what happens when you guys know this information ahead of time? I'd be curious about that.

For those who did the beer game, one of the things that came out of that that's interesting, the beer game is this multi-echelon supply chain. The ultimate thing about it is that information matters, information is helpful. When you play the beer game with people who have played the beer game, you still get the exact same outcome no matter how many times you do it.

And one of the classic examples is we had a beer game, like, championship here at MIT some years ago where not only was it people who played before but it was professors who teach the beer game in the same room. And we also gave information. So typically in the beer game, there's some--

AUDIENCE: [INAUDIBLE]

JAMES PAINE: I know, yeah. [INAUDIBLE] But even the folks who teach the beer game didn't escape the beer game. So part of this is that as long as-- and this is the whole idea of structure yields behavior. The structure is still there. The underlying structure still exists. It is real hard to break your behavior patterns when you don't change that underlying structure.

So if you want more, we're all up on the fourth floor, for one thing. So at least, I don't lock my door, so come on in. I'll talk your ear off.

For classes here at MIT, a lot of what we just talked about is covered really in much more detail in 15.871, Introduction to System Dynamics. I would say 15.871's primary purpose is to get these ideas of systems thinking in your head. You'll be spending a lot of time in that class up on a chalkboard, sort of talking through problems, saying, OK, well, let's talk about what happens when you increase consumer satisfaction. What does that mean?

And start drawing these lovely causal loop diagrams and think to ourselves, OK, well, what happens if we snip this loop here? And what would it take to do that? So it's much more about how to get yourself thinking from linear direction into more of a looping direction.

15.872 dives a lot more into now taking those thought processes and applying them to sort of much more realistic scenarios and situations and getting more comfortable with modeling tools. Both involve you using software sets, which I'll talk about a little bit later, that are commonly used within system dynamics. But 15.872 kind of takes the training wheels off. And there is no-- there is no right there is no wrong answer as long as you can figure out a good self-contained simulation that gets the conclusion you think is interesting.

This one is relatively new, 15.873. This is somewhere that sits in between 15.871 and 872. This is one who's, I like to think, has more of an emphasis on business and policy.

I have this example down here. This is the Vasa, which I'm now giving away a slide from, I think, 15.872. The Vasa was an incredible ship that was built by some of the best ship designers of its day, was launched, and then promptly tipped completely upside down and sunk to the bottom of the ocean.

And the reason was because, at the last minute, there were several late design adds. Someone decided they wanted to build the captain's lodge a little bit larger. They wanted to add a few more sails. They wanted to add a ton more cannons. They wanted to make it more and more impressive. It became hydro dynamically unstable, promptly tipped over.

So the idea being here that-- I emphasize that one for this idea that you have unintended consequences. There are no side effects. There are simply effects that you didn't plan for. And that's kind of emphasized a little bit more in 15.872, 15.873.

In terms of books, I brought a few with me. Again, I mentioned this at the beginning of the class. This has become sort of the bible of system dynamics, business dynamics. The first couple of chapters cover a lot of what we talked about.

Later on, it gets really specific. There's an opening in here that says, essentially, if you have a background in math, that's great. If you don't, that's great. But later on, it starts talking about the specific differences between choosing modes of feedback and behavior, how to model specific things, really how to get your hands dirty for primarily compartmental models.

If you want something a little more general, there is the Fifth Discipline Field Book. I wouldn't necessarily recommend the whole book unless that's your thing. But part two is called systems thinking. It does a great job of summarizing the work that we talked about here.

Both of these books are in the Dewey Library. I checked earlier today. I think this is one of the two copies of this book. Sell or return it right after this class. But there's lots of copies of John's book on business dynamics.

If you want to know a little bit more about sort of the history of this, Limits to Growth is a book that I recommend because it has a strong history in system dynamics. It also has a bit of a controversial history, if you dive into it. This was one of those first books that applied the ideas of overshoot and collapse to the earth, and saying that we as a species are headed towards this, and we need to manage our resources and we need to be a little bit more careful with our choices.

The simulations in there had some specific dates that ultimately were held up as sort of the how why point prediction is difficult in system dynamics. So if you read it, this is not a prediction for the future. It's discussing the mode of behavior. It's talking about the structure that we currently have that can lead to overshoot and collapse within shared resources.

Articles, if you get a chance, it's my personal favorites. System Dynamics at Sixty: The Path Forward, I think, does a great job of kind of saying where is system dynamics right now. A lot of material that you get out there talks about where system dynamics was 20, 30, even 40 years ago. I think this one does a great job of saying, that's great. What are we doing from now?

This is essentially the article that says, system dynamics is not compartmental models. System dynamics is a mode of thinking. Choose the software that you need to. Choose the methods of communication you need to.

I am, again, someone who does a lot of operations management. So most of these choices up there about compatibility traps and sort of operational failure, that's kind of where I'm coming from. But that being said, this one right here, the third one on the list, Making the Numbers, applies a classic operations management concept of a capability trap to valuation of stocks. And I think that's a great little connection right there. If you happen to be on MIT'S network, these are all free and accessible. The worst case scenario, I'll print out a copy for you because I think these are good articles.

Other websites. Some personal favorites here. Creative Learning Exchange, this is one-- it's primarily K through 12 education-focused. But there's a lot of fun little tools on there.

It uses a lot of animations, make things very approachable, especially if you're used to looking at sort of cold, black and white diagrams. Clicking around on there just kind of makes me happy, like it has lots of pink fluid flowing back and forth, and you like-- you attach the little things together. They have some fun iOS and Android applications.

Tom Fiddaman's MetaSD website is fantastic. He has-- that man will just model anything that he thinks is interesting. So if you have a blank piece of software in front of you and you don't know what to click and you want something as an example, he has one about sort of the self-reinforcing nature of UFO sightings.

So he essentially thinks to himself, this is an interesting behavior I have observed. I wonder what structure could have yielded it? And then he makes it and puts it on his website. So that's just a fun one to click around on.

This one, the self-study website, that is dense, to be honest. But it's a good spot to start. This systematically walks you through the fundamentals of system dynamics from open concept to detailed modeling systematically. And is free and open for anyone to use.

I run, right now, the Journal Club. So that's my plug right there. Stop on by. Every other Friday, once the semester starts, or every Friday once the semester starts, we have a seminar series that's systematic-focused up on the fourth floor. If there's not a seminar scheduled, then I just kind of run the discussion and talk about something I feel like talking about.

And then of course, the System Dynamics Society. That's kind of our group society. I would be remiss if I did not at least mention them in passing. It's a lot of information, but I think especially those first two are not ones that you see a lot are recommended.

Software. We didn't talk at all about software. If this had been a longer course, that probably would be the next thing I'd do is I'd dive into some software with you guys. Unfortunately, we don't have the time.

So I want to point you in these-- towards these guys right here, Vensim and Stella Architect. Vensim especially has a free academic license. If you also have an MIT email address, you can get more or less the full commercial version for free, which is nice. The difference between the personal learning edition and the commercial edition is minimal until you hit it. And then once you hit it, you really want that commercial version for being able to just do some stuff. So if you have an MIT email address, I suggest spending the three minutes to get a Vensim professional license.

Stella Architect is less used in our group, but definitely used in general. They have a great storytelling mode, this idea of building up your models progressively without scaring people away, kind of. And also have a lot of nice visuals. I'm still getting used to that software myself, but I like what I see.

Also, off to the side there, NetLogo. Not used a lot in our group necessarily, but I use it a lot. It's a good old agent-based model.

So sometimes it's useful to think, OK, I know what a person would do or an agent or an entity in the setup. I can wrap my head around that. I don't know what the system is, but I can wrap my head around one person wandering through space. This lets you, essentially, model that person and then put a whole bunch of them in a room and see what happens. If anyone has done any sort of agent-based modeling with any sort of scripting language, MATLAB, R, anything else like that, this is kind of the same concept with a nice little pretty wrapper around the outside of it.

Speaking of which, in terms of software, we use all sorts of software. It's not just Vensim and Stella. These are just what I have upstairs on my machine right now. Actually, they're on my machine right now and what I use on more or less a daily basis.

We're not just compartmentalize models. We're a mode of thinking. So if you have a tool that gets the job done, than good for you. It's possible to do a good causal diagram using Excel if you're really motivated.

Oh yeah, and a whiteboard. Whiteboard definitely gets the job done. That's actually-- he raises a great point. The best use is just standing in front of our whiteboard and starting writing some things down.

People often ask the question, what do you do? Like, how do you start? Like, what's the first thing you do?

I'd say walk up to a whiteboard, think about the thing you care about. I don't know, it's CO2 in the atmosphere, car deaths, number of people quitting from a non-profit. Write that on a board. Just write it smack dab in the middle of the board.

Then think to yourself, what would cause that thing to go up or down? Write it off to the side. Draw an arrow, put a plus or a minus next to it. And just keep going and see what happens. And let those loops come back around. And that's a good spot to start.

So that's-- we have a few minutes. I'm going to click this one thing right here for another example of system dynamics in action. If anyone is familiar with the policy interactive, this is an idea too of what's the most complicated model sort of idea and what can you do.

This is En-ROADS. This is a policy simulator for the climate. So this right here has a full-blown climate simulation running on the back. But its job at the end of the day is again, not point prediction. Its job is policy.

So again, this is one kind of a fun thing to click around and kind of see system dynamics in action. This is simply saying, what are we doing right here? What's our predicted temperature change over time? What are the choices that we can make as a society, as a group, as a series of countries to change this outcome?

This is being used right now in conversations with members of our own government in order to help elicit mental models and get people on the same page about climate action in the future. And so this is something else where there is a big compartmental model running in the background. But this is an example of using system dynamics to hopefully elicit real policy change in the near term. So you can start clicking around, seeing what happens.

I'll be honest, I'm really bad at this one. I click around and I realize I'm really bad at figuring out how to fix the climate, which is why I do operations and supply chain management, I suppose, at the end of the day. So yeah, I'd encourage you guys to click around on that.

So that, essentially, is where we are. Thank you guys very much. I loved sharing system dynamics with you, and I hope to see some of you guys in the future. Thank you.

[APPLAUSE]