TIME and CAUSALITY IN THE SCIENCES

June 7-9, 2017
Hoboken, NJ

STEVENS
INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®

Photo by Dmitry Avdeev

Welcome

Welcome to the 12th conference in the Causality in the Sciences conference series. Each year these conferences bring together philosophers and scientists to explore various aspects of causality. This edition focuses on the relationship between time and causality across philosophy, computation, and specific scientific disciplines.

Time is central to how we perceive and learn about causes, and a key aspect of many definitions of causality is that the cause must occur before the effect. Yet, time can also be misleading, as one event following right after another may give the impression of causality where none exists, and events may be observed out of order because of variables having different sampling rates. Further, as we collect increasingly large longitudinal datasets, we need new ways to efficiently extract causal relationships, and the methods must be robust to missing variables, noise, and causal relationships that change over time.

This conference aims to foster research of these and other facets of the relationship between time and causality in the sciences. We hope the range of interdisciplinary perspectives will lead to thought-provoking discussions.

With warm wishes on behalf of the organizers and CitS steering committee, Samantha Kleinberg

Organizers

Samantha Kleinberg, Stevens Institute of Technology

Michael Strevens, New York University

Causality in the Sciences Steering Committee

Phyllis Illari, University College London

Bert Leuridan, University of Antwerp

Julian Reiss, Durham University

Federica Russo, University of Amsterdam

Erik Weber, Ghent University

Jon Williamson, University of Kent

Logistical Information

Wireless Internet Access

Please use the information below to access the wifi connection:

Network:	Stevens Guest
Username:	tacits17
Password:	yellow6bird

Campus map

Key locations

Conference sessions take place in Babbio 122. Coffee breaks and lunches are located in the Babbio atrium, and the reception on June 8 will take place on the 5th floor patio in the Babbio center (weather permitting).

Dinner cruise information

Transportation to Dinner Cruise

5:00 p.m. End of Conference Day 1
5:45 p.m. Transportation from Babbio Center to marina
6:00 p.m. Arrive at Lincoln Harbor Yacht Club Marina and board yacht
6:30 p.m. Cruise departs

Transportation from Dinner Cruise to Stevens Institute of Technology

9:30 p.m. Cruise Returns to marina
9:45 p.m. Transportation from marina to Babbio Center

Pick-up Location

Transportation will be located outside the main entrance of the Babbio Center on River Street.

Babbio Center - Stevens Institute of Technology

525 River Street Hoboken, New Jersey

Lincoln Harbor Yacht Club Marina

1500 Harbor Blvd. Weehawken, New Jersey Complimentary parking is available across the street from the yacht boarding area. It is a self-park lot located directly across from 1500 Harbor Blvd. and in between the Houlihan's and Estuary apartment complex.

Schedule

Wednesday, June 7

8:30-9:15	Breakfast and registration
9:15-9:30	Welcome
9:30-10:30	Keynote: **Betsy Ogburn** Social networks, causal inference, and chain graphs
10:30-11:00	Coffee break
11:00-12:30	Philosophy of Social Science: **Jonathan Livengood and Karen Zwier** Temporalization in Causal Modeling **Julian Reiss** Time Series, Non-Stationarity, and Causal Inference **Yin Chung Au** Registering time-dependent events for searching cell mechanisms
12:30-1:30	Lunch
1:30-3:00	Psychology and Medicine: **Neil Bramley, Ralf Mayrhofer, Tobias Gerstenberg, and David Lagnado** Causal learning from interventions and dynamics in continuous time **Tobias Gerstenberg, Noah Goodman, David Lagnado, and Joshua Tenenbaum** A counterfactual simulation model of causal judgment **Michael Wilde** Another disambiguation of the Russo-Williamson Thesis
3:00-3:30	Coffee break
3:30-4:30	Keynote: **Marc Buehner** Perception of Time and Judgment of Causality: Mutual Constraints
4:30-5:00	Epistemology **Jürgen Landes, Barbara Osimani, and Roland Poellinger** Probabilistic Causal Inference through Evidence Synthesis
5:00-6:00	Break and transport to cruise
6:00-9:30	Boat cruise and dinner

Thursday, June 8

9:00-9:30	Breakfast
9:30-10:30	Keynote: **David Jensen** Overcoming the poverty of mechanism in causal models
10:30-11:00	Coffee break
11:00-12:30	Computation: **Daniel Malinsky** Causal inference from time series data with unmeasured confounding: how data-driven can we get? **Yi Zhao and Xi Luo** Granger Mediation Analysis of Functional Magnetic Resonance Imaging Time Series **Naftali Weinberger** Can Causal Models Include Variables with Their Time-Derivatives?
12:30-1:30	Lunch
1:30-3:00	Philosophy of Physics: **Inge De Bal and Erik Weber** Mechanistic versus correlational evidence in reasoning about physical claims **Toby Friend** Could laws determine determinism? **Keming Chen** Our Knowledge about the Past: Some Puzzles about the Past Hypothesis, the Principle of Indifference, and Self-Locating Uncertainty
3:00-3:30	Coffee break
3:30-4:30	Keynote: **Jenann Ismael** Time and Causation: New Insights on Old Topics
4:30-5:00	History of Philosophy of Science: **Flavia Padovani** At the Roots of the Common Cause
5:00-6:00	Reception at Babbio center, 5th floor patio (atrium in case of rain)

Friday, June 9

9:00-9:30	Breakfast
9:30-10:30	Metaphysics: **Graeme A Forbes and Veli-Pekka Parkkinen** Delayed Causation **Victor Gjisbers** On the Causal Nature of Time
10:30-11:00	Coffee break
11:00-11:30	Philosophy of Science: **Thomas Lodewyckx and Bert Leuridan** Causation and Time: Synchronic Causality
11:30-12:30	Keynote: **Phil Dowe** The Direction of Causation
12:30	Closing remarks

Keynote abstracts

Marc Buehner

Cardiff University

Perception of Time and Judgment of Causality: Mutual Constraints

In this talk I will explore how perception of time influences causal judgment, and, how in turn causal knowledge influences temporal experience. I will review previous research on the role of temporal contiguity in causal judgments, and will outline that temporal regularity (or predictability) is a further important cue to causal judgments. Most standard theories of causal learning (whether based on associative or rule-based learning) cannot easily represent this role of regularity, but prior-knowledge driven / evidence integration accounts (e.g. Bayes) can. Bayesian accounts also fit well with the second aspect of the talk — systematic distortions of time perception in the presence of causal knowledge. There is now a substantial body of research showing that time perception is malleable by context, such that the same objective interval is perceived differently when the events demarcating it are linked by a causal connection. Specifically, causal intervals are perceived as shorter than non-causal intervals of identical length, and causes and their effects mutually attract each other in subjective space-time. The overall pattern of evidence – that perception of time and judgment of causality mutually constrain each other – fits well with cognitive theories that assign a critical role to causality.

Phil Dowe

Australian National University

The Direction of Causation

Two events a, b may be causally connected in the sense that either a causes b or b causes a. What makes it so that a causes b rather than b causes a (or vice versa)? Various answers have been given to this question, including temporal order, agency, certain patterns of correlations, and primitive causing. It may seem remarkable that plausible cases can be made for such different answers. I show why none of these is completely satisfactory from a metaphysical point of view, and set out in general terms what it would take to give a satisfactory answer.

Jenann Ismael

University of Arizona

Time and Causation: New Insights on Old Topics

Philosophical inquiry into the nature and time and causation has been around as long as there has been philosophy. New insights have emerged in the last twenty or thirty years on both topics, which developed out of contributions from physics, computer science, and decision theory. I want to pull those new insights together, convey the picture that is emerging, and see how it answers some traditional questions. I will end with a discussion of new questions raised by the emerging picture.

David Jensen

University of Massachusetts, Amherst

Overcoming the poverty of mechanism in causal models

Work over the past several decades on both causal graphical models and the potential outcome framework has yielded significant advances in our understanding of causal inference. However, a substantial gap remains between the types of causal models expressible within these frameworks and what practicing scientists view as adequate causal explanation. In this talk, I will survey the extent to which these formal frameworks can represent and reason about causal mechanism, describe various recent developments that move toward richer formal descriptions of mechanism, and suggest research directions that could provide a greater alignment between formal accounts of causal reasoning and traditional scientific theories.

Betsy Ogburn

Johns Hopkins University

Social networks, causal inference, and chain graphs

Traditionally, statistical and causal inference on human subjects relies on the assumption that individuals are independently affected by treatments or exposures. However, recently there has been increasing interest in settings, such as social networks, where treatments may spill over from the treated individual to his or her social contacts and where outcomes may be contagious. Researchers interested in causal inference have developed methods for interference – when one individual's treatment or exposure affects not only his/her own outcome but also the outcomes of his/her contacts– and researchers interested in social networks have attempted to model the spread of contagious outcomes across network ties. In both of these settings, causal inference using non-experimental data requires observing longitudinal data on treatments and outcomes as they evolve in real-time, so that each spillover or contagious event appears in the data. This results in two roadblocks for researchers. First, in most settings it is impossible to collect the kind of real-time data required. The time intervals for longitudinal data collection must be short enough to capture every potential transmission event, which could mean weeks, days, or even minutes or seconds. Second, even if the full longitudinal data are available, the resulting model will generally be high-dimensional and often too big to fit to the available data. As a practical matter, most researchers deal with reduced data, comprised of observations collected at one or a small number of time points. We propose and justify a parsimonious parameterization for social network data with interference and contagion. Our parameterization corresponds to a particular family of graphical models known as chain graphs. We demonstrate that chain graph models approximate the projection of the full longitudinal data onto the observed data, which is missing most of the time points from the full data. We illustrate the use of chain graphs for approximate causal inference about contagion, interference, and collective decision making in social networks when the longitudinal evolution of treatments/outcomes is not fully observed.

Contributed abstracts

Listed alphabetically by last name of first author.

Yin Chung Au

Registering time-dependent events for searching cell mechanisms

This case study of cell biological mechanism research argues that registration and representation of both time and time-dependent events play crucial roles in the perception of causalities within the targeted mechanisms. Following Glennan's distinction between two kinds of causal relation at different levels of natural selection, here I view causality in cell mechanisms as referring to either causal relevance or causal productivity. Importantly, perceptions of both kinds of causal relation can contribute to constructing the same 'pathways'.

In the practice of searching cell mechanisms, researchers may first address the temporal sequence of events observed and then may or may not address the causality. This depends on whether they obtain sufficient observation of relevant events between the two ends of the sequence. Normally, given the complexity of biological mechanisms, the causality in turn comprises of underlying mechanisms. In controlled experiments, observation of how intervention changes the sequence of events is key to inferring causality and revealing underlying mechanisms.

This study extends the study on mechanism diagrams of circadian rhythms by Bechtel et al. (2015) by showing that recognition of, and emphasis upon time in biological mechanisms are not limited to research that is obviously about time (circadian pattern). This study goes further to show that, in biological experiments for searching mechanisms, the design for a time frame of intervention and observation can sometimes be arbitrary and based on existing knowledge of comparable conditions and/or properties of similar entities. Thus the registration of time and time-dependent events becomes more important for constructing the pathway via which a sequence of events occurs.

In most instances, when drawing cell mechanism diagrams, researchers remove the visual representations of time, such as time axis and symbols of time courses, i.e. what Bechtel at al. call the spatial dimension(s) used to represent time. In some other instances, nonetheless, researchers retain the representations of time and even make them more explicit by visually interweaving them and the observations. In this case, the time course of experimental design tends to be visually emphasised. This study uses examples of both cases (time is absent and time is present in mechanism diagrams) to show that, in both cases, there are intercalations of time prediction and observation into each other during the experiment. Time prediction is derived from existing 'store' (borrowed

from Craver and Darden), and observation is for the time-dependence of events induced by intervention.

The sampling and selection of biological papers for this study are based on a database composed of research and review papers in the apoptosis (programmed cell death) field from the 1970s to present. This study analyses the whole paper, including text, data images and diagrams, in order to find out the relationship between perception of causality and registration of time/ time-dependent events in the development of biological arguments.

Neil Bramley, Ralf Mayrhofer, Tobias Gerstenberg and David Lagnado

Causal learning from interventions and dynamics in continuous time

Event timing and interventions (actions that manipulate causal variables) are important and closely related cues to causal structure (Lagnado & Sloman, 2004). However, they have typically been studied separately. Here, we bring them together for the first time both empirically and theoretically. We present an experiment in which participants learn causal structure through free selection of interventions on causal devices whose dynamics unfold in continuous time. In parallel, we develop a Bayesian model that infers structure from event dynamics. We contrast learning in devices with and without feedback loops, and where the true cause–effect delays are either more or less reliable.

We find that successful learners use their interventions to structure and simplify their interactions with the devices. Participants typically spread their tests out evenly in time in ways that minimise causal ambiguity and minimise the costs of inference. Crucially, this tendency was unrelated to the information produced, from ideal learning perspective, and is we propose that it is indicative of a cognitive strategy for structuring and simplifying learning.

Consistent with work in causal learning from contingency information (e.g., Bramley, Lagnado & Speekenbrink, 2015; Coenen, Rehder & Gureckis, 2015), we find that participants form a non-normative preference for intervening on putative root components. Additionally, we find much lower accuracy for inferring the structure of devices containing feedback loops. Building on this, we identify particular loci of evidential complexity and resulting inferential difficulty resulting from the presence of a cycle in a larger causal system.

Finally, digging deeper into participants' judgment patterns and deviations from normativity, we also consider a range of learning heuristics based in limited timing sensitivity and reliance on online and local inference. We find we can capture participants best with a model that incrementally constructs, tests and adapts a single structural hypothesis, consistent with recent work on com-

plex causal structure induction from contingency data (e.g., Bramley, Dayan, Griffiths & Lagnado, 2017).

References

Bramley, N. R., Mayrhofer, R., Gerstenberg, T. & Lagnado, D. A. (submitted). Causal learning from interventions and dynamics in continuous time.

Bramley, N. R., Dayan, P., Griffiths, T. L. & Lagnado, D. A. (2017). Formalizing Neurath's ship: Approximate algorithms for online causal learning. Psychological Review, to appear.

Bramley, N. R., Lagnado, D. A. & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through interventions. Journal of Experimental Psychology: Learning, Memory & Cognition, Vol 41(3), 708-731.

Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. Cognitive psychology, 79, 102-133.

Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. Journal of Experimental Psychology: Learning, Memory, and Cognition, 30(4), 856.

Eddy Keming Chen

Our Knowledge about the Past: Some Puzzles about the Past Hypothesis, the Principle of Indifference, and Self-Locating Uncertainty

In statistical physics and philosophy of physics, it has been standard to follow Ludwig Boltzmann and impose a Past Hypothesis on the boundary of the physical space-time. According to the Past Hypothesis, the "initial" state of the universe is in a very orderly (low-entropy) state. In this talk, I would like to explore an alternative hypothesis, motivated by the (in)famous Principle of Indifference. I will argue that both theories, on pain of empirical inadequacy, require certain objective self-locating prior distributions. It follows from reasonable premises that the two theories are in fact empirically adequate to the same extent. This is a surprising result, for it leads to deep and puzzling consequences for the epistemic justification for our beliefs about the past (including the prosaic ones that we and our surroundings were "younger" in the past). We will then think about what this might mean for the general issues in philosophy of science about theory choice and pragmatic considerations.

Inge de Bal and Erik Weber

Mechanistic versus correlational evidence in reasoning about physical causal claims

In this paper we investigate the interplay between mechanistic evidence and correlational evidence (as encoded in the laws of physics) in reasoning about physical causal claims. Consider two examples of physical causal claims:
(a) The height of the flagpole and the position of the sun determine the length of the shadow of the flagpole.
(b) In pressure cookers, the temperature of the gas determines the pressure it exerts on the walls of the cooker.
The decision to accept these claims (and reject opposite claims, e.g. the height of the flagpole and the length of the shadow determine the position of the sun) cannot be based solely on physical laws, because almost all these laws are symmetrical (temporally, but also with respect to the inferences they support).

The interaction between flagpole, sun and shadow, for example, is expressed in the following law of geometrical optics:

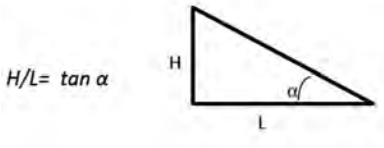

H is the height of the flagpole, L the length of the shadow and α the angle of elevation of the sun above the horizon. Using this equation, we can calculate the value of each variable based on the other two. In this way, we can construct an equally strong argument for causal claim (a) or any of its converses (which, if true, would involve some kind of backward-in-time causation). In order to decide in favour of (a), we need more information. There is an argumentative gap between the symmetric physical laws and the asymmetric physical causal claims.

Our aim is double:
(1) To show that it is mechanistic evidence (knowledge about the underlying mechanism) that fills this argumentative gap.
(2) To explicate how the two kinds of evidence interact with each other so that they can provide good reasons for accepting physical causal claims.

In order to reach these aims, we draw on an analogy with an evidential condition that is common in the social sciences. This condition can be illustrated by means of the evidence for a causal claim in political science known as Duverger's law:

> The simple-majority single-ballot system favours the two-party system.

In a simple-majority single-ballot system there is one member of parliament to be elected in each voting district. The candidate who gets more votes than any other candidate is elected (even if there is no majority, i.e. the candidate's score is less than 50%). Duverger considers two other systems: the majority system with a second-round runoff (if no candidate receives more than 50% of the initial votes, there is a second round with the top-two candidates) and proportional representation (multiple members of parliament for each district; seats allocated based on percentage of votes for each political party). Duverger's causal claims about these systems are:

> The majority system with a second-round runoff favours multi-partism.
> Proportional representation favours multi-partism.

The evidence for Duverger's three claims is partly correlational (frequency of co-occurrence between types of electoral systems and number of parties) and partly mechanistic (a psychological mechanism and a mechanism for translating votes into parliament seats). Correlations are symmetric, the mechanisms are necessary to support conclusions about (temporally asymmetric) causal claims.

We analyse how the two types of evidence interact in this case and in related cases in the social sciences. Then we argue that the way in which mechanistic evidence and physical laws interact is similar. We also point at an important difference regarding the level of generality of causal claims.

Graeme A Forbes and Veli-Pekka Parkkinen

Delayed Causation

There has been considerable attention paid to action at a spatial distance, with work on things like gravitons, and fields of force being particularly significant. Less attention has been paid to action at a temporal distance. There has been work on *backwards* causation, where the cause comes after the effect, but the phenomenon where the effect follows the cause by some considerable distance of time has been neglected. It seems, at least, that some form of delayed causation occurs commonly. The budget deficit is attributed to the previous government many years after they have left power, and the cavities in one's teeth are attributed to omitting to brush them properly as a youth. What is delayed causation, what varieties does it occur in, and how can one intervene upon it? We will develop a brief taxonomy, distinguishing between chains of similar causes the effects of which happen without considerable delay, but only become cumulatively noticeable after a long period of time, and cases where the effect is the product of some mechanism that is only triggered after some threshold is passed. In addition to this taxonomic work, we will defend one negative claim. The mere passage of time cannot be a cause – there must always be some combination of chains of causation and triggered mechanisms to explain any instance of delayed causation. These are, broadly speaking, metaphysical assumptions based on intuitions about causation. However, this is not just a matter of metaphysics. A closely similar assumption is involved in causal modelling frameworks based on the Causal Markov Condition, which implies that the direct causes of an effect screen it off from its more distal causes. Applying this idea to causal chains where directness is understood as temporal, it implies that the temporally most proximate causes must screen the effect off from its temporally more distant causes. There are cases where this condition is seemingly violated. We will investigate two such examples, one a general type of causal process and the other a particular case: processes of ecological succession, and the effect of the Exxon Valdez oil spill on the fisheries of Prince William Sound. We argue that in such cases, a distant cause is structure-altering in that it changes the mechanisms by which the effect is normally produced. That such cases violate the Markov condition is a symptom of the fact that the models to which the condition is applied do not represent the mechanisms responsible for the conditional probabilistic dependencies between variables. We conclude that the fact that the condition is violated in such cases does not imply that we should give up the metaphysical considerations that speak against delayed causation in general.

Toby Friend

Could laws determine determinism?

Despite plenty of variation in the definition of cross-temporal causal determinism, it is typical to consider it a matter of how the laws are: the laws determine whether determinism of causation exists from one time to another (e.g. van Inwagen 1975, Earman 1986). It is my aim to show that it is very difficult to make sense of this idea and thereby put into question whether we have ever had good reason to be determinists. In contrast with existing literature, my aim is not to show that specific theories are indeterministic (e.g. Earman 1986, Norton 2008). Rather, my claim is more general: laws are just the wrong sort of thing to determine whether determinism holds. Moreover, despite the widely acknowledged indeterminism of contemporary physics (not to mention in other sciences) I believe my claim is of more than mere academic interest. Although I will not pursue the idea in my presentation, the argument I provide suggests that the general idea of laws licensing any kind of inference across time is also highly problematic.

My argument begins with the commonplace schema,

$$\forall x(Fx \rightarrow Gx).$$

If this is a general schema for laws we see immediately why laws could never be sufficient for determinism. The absence of a time-variable indicates that such propositions describe a simultaneous relationship between property-instances. Hence, they cannot support a temporally extended causal implication. But I will further show that it is of no help to advocate an alternative schema in which a temporal duration exists between the property instances, i.e.,

$$\forall x \forall t(F(x,t) \rightarrow \exists y(G(y,t+\epsilon)))$$

(for times t and duration ϵ). For the retained conditional form means that a conjunction of initial conditions and laws can only determine a future state under the additional assumption that the initial conditions all fall under some laws' antecedent predicate. In effect, what is needed is a further law of the form,

$$\forall x(Fx).$$

But I argue that there are no such laws. According to Author (2016), all laws predicate behaviour *if certain system-conditions are satisfied*. Hence, the conditional form is unavoidable. Moreover, this relationship is always synchronous. I will show how this reasoning can be applied even in the case of laws which relate variables at different times (e.g. growth models) and in the case of our most general laws (e.g. the wave-equation, the total force law, conservation principles).

A different tack one might pursue in order to reestablish the coherency of a law-determined determinism invokes meta-laws (e.g. symmetry principles).

These make claims about what the laws must look like, thereby restricting the possibility of undetermined but lawlike interactions. However, I suggest that this suffers the same issues as we have just seen but at the meta-level. For meta-laws plausibly have their own meta-conditions, i.e. they will not make claims about all laws in general but only about particular classes of laws. Hence, no conjunction of meta-laws will be alone sufficient to preclude a law which is not an instance of a meta-law. Regardless, even if there are meta-laws concerning all laws, they alone would not imply that all worldly behaviour is law-governed.

To conclude, I suggest that in order for any cross-temporal causal determinism to exist, it needs to be the case not only that laws of a certain kind are true but that the highly contentious claim that the laws exhaust all possible phenomena is true. This should lead us to question both why it was ever assumed that the laws were sufficient for determinism and, derivatively, why determinism was ever a live option at all.

References

Earman, J.: 1986, *A Primer on Determinism*, D. Reidel Publishing Company.

Norton, J.: 2008, The Dome: An Unexpectedly Simple Failure of Determinism, *Philosophy of Science* 75, 786–798.

van Inwagen, P.: 1975, The Incompatibility of Free Will and Determinism, *Philosophical Studies* 27(3), 185–199.

Tobias Gerstenberg, Noah Goodman, David Lagnado and Joshua Tenenbaum

A counterfactual simulation model of causal judgments

How do we make causal judgments? In this talk, I will present a counterfactual simulation model (CSM) of causal judgments that combines key insights from process and dependence theories of causation. The CSM predicts that people's causal judgments are influenced by the extent to which a candidate cause made a difference to i) whether the outcome occurred, and ii) how it occurred. I will show how whether-causation and how-causation can be expressed in terms of different counterfactual contrasts defined over the same generative model of a domain. I will focus on applying the CSM to the domain of intuitive physics, asking people to make judgments about colliding billiard balls. The CSM accounts for participants' causal judgments to a high degree of quantitative accuracy. Causal judgments increased the more certain participants were that a ball was a whether-cause, a how-cause, as well as sufficient for bringing about the outcome.

The CSM postulates that people make causal judgments by comparing what actually happened with what would have happened if the candidate cause had been removed from the scene. In direct support of this claim, I will show eye-tracking data of how people mentally simulate how the relevant counterfactual situation would have unfolded. To demonstrate the flexibility of our account, I will show how the CSM naturally captures causal judgments about omissions, judgments of causal responsibility for the stability of towers, and can help us better understand the mapping between causal events in the world and the words we use to describe them.

While the model does a good job of capturing people's judgments overall, there are still some situations it struggles with. Consider the case of preempted causation: Ball A knocks ball E into the goal before ball B would have done the same. Clearly, ball A caused ball E to go in the goal (and not ball B). However, the fact that the presence of ball B makes no difference to participants' judgments (about ball A) is somewhat puzzling. Shouldn't the fact that B would have caused the same outcome to come about reduce causal judgments to ball A? Interestingly, the pattern of results looks very different for preempted prevention. Consider a situation in which ball E is headed toward the goal and ball A knocks ball E out of the way shortly before ball B would have done the same. Here, participants are less willing to say that A prevented E from going into the goal when ball B would have done the same just a moment later. Together, this pattern of findings illustrates an important asymmetry in the role that temporal information plays for judgments of causation and prevention. I will present some ideas for how to resolve this puzzle by using temporal information to construct a causal representation of the situation which constrains what counterfactual contrasts are considered.

Victor Gjisbers

On the Causal Nature of Time

Most contemporary philosophers believe that time and space are conceptually prior to causation: we may have to analyse causation at least partly in terms of space and time, e.g., in terms of the spatiotemporal locations of particular events, but we do not have to analyse space and time in terms of causation. In one way or another, this idea is assumed by almost all contemporary theories of causation. But the thesis is rarely defended or even made explicit. In this presentation, I will focus on the question whether time is conceptually prior to causation and I will develop an argument to the contrary.

Perhaps the last influential attack on the priority of time over causation was the development of causal theories of time in the 50's and 60's by authors such as Reichenbach, Grünbaum and Van Fraassen. These authors aimed to reduce the concept of time to the concept of causation, as Leibniz had tried to do several centuries earlier. Their efforts were forcefully attacked in the late 60's and early 70's by authors such as Lacey and Earman, leading to an abandonment of the project.

Rather than resurrecting this reductionist approach, I wish to develop a line of thought that harkens back not to Leibniz, but to Kant; more specifically, to his Second Analogy of Experience. In the Second Analogy, Kant argues that all events have to be in time; but that an event can only have a determinate place in time if it is linked to other events through an exceptionless causal law.

While Kant's argument may suffer from its reliance on execeptionless laws and the metaphysics of transcendental idealism, an updated epistemological version of it can be developed. The first stage of my updated argument consists of the claim we have to choose between two horns of a dilemma: either we adopt as an a priori principle the claim that the causal order is identical to the temporal order – that is, that causes are always earlier in time than their effects – or we have to embrace complete scepticism with regards to the temporal location of any and all events. Since extreme scepticism is unpalatable, we have to choose the first horn.

In the second stage of the argument, I show that the status of the a priori principle cannot be that of a conventional definition of cause and effect; if it were, it could never do the epistemological work that it has to perform. Instead, the principle expresses a necessary link between time and causation, a link that is so strong that we can't really be said to have a concept of time unless we also have the concept of causation.

If successful, this argument should change our approach to theories of causation, as well as leading us to abandon the influential Lewisian idea that there are possible worlds without causal regularities.

Jürgen Landes, Barbara Osimani, and Roland Poellinger

Probabilistic Causal Inference through Evidence Synthesis

Current methods for the purpose of causal inference aim to deliver a categorical assessment about the presence of a causal relationship between events or variables. This is at odds with the great amount of epistemic and ethical uncertainty surrounding most applied sciences. In particular, for the sake of the precautionary principle, this uncertainty should not be dismissed but rather explicitly accounted for in detecting, preventing and managing e.g. environmental or health hazards (Kreibel et al. 2001, Raffensperger, and Tickner 1999). The rationale for the attenuation of the requirement of scientific proof and certainty about the causal link is ultimately one of minimising expected loss by anticipating risk detection and prevention (Osimani & Russo 2016; Osimani, Russo, Williamson, 2011).

We here present a framework for causal assessment which allows the incorporation of heterogeneous pieces of evidence via a probabilistic judgement about the causal link between candidate causes and effects (Landes, Osimani, Poellinger, 2017). The framework comes in the form of a Bayesian network whose nodes represent epistemic variables related to causal associations. In particular, our system i) identifies possible indicators of causality on the basis of the methodological and philosophical literature on causality, evidence, and causal inference; ii) embeds them in a topological framework of probabilistic dependencies and independencies grounded in assumptions regarding their reciprocal epistemic interconnections; iii) weakly orders some of these probabilistic dependencies as a function of their inferential strength with respect to the confirmation of causal hypotheses. This system has been developed for the purpose of drug safety assessment, but it can be easily applied to other domains with relatively few adjustments.

Our framework accommodates a number of intuitions already expressed in the literature concerning the EBM vs. pluralist debate on causal inference, evidence hierarchies, causal holism, relevance (external validity), and reliability (see for instance, Howick 2011, Clarke et al. 2014, Cartwright 2011, Teira 2011). In the talk we will particularly focus on precedence in time as an (imperfect) indicator of causality in relationship to other abstract indicators and to its concrete role in causal inference from observational data of various kinds. This role may be very complex when causal relationships are not conceived as instantaneous but rather lasting in time, or proceeding from cumulative incrementation of the causal force (for instance through incremental dosage, or simply through prolonged risk exposure over time). Frequency of effects in time can also be a deceiving signal, especially if exploited by interested parties who succeed in manipulating it (as for instance in the Vioxx case). This will give us the opportunity to illustrate the social epistemology background in which

causal inference is embedded and to show how our approach can take also this dimension into account and incorporate it in the final causal judgment.

Jonathan Livengood and Karen Zwier

Temporalization in Causal Modeling

Causal influence, as it is modeled and discussed in the social sciences, is widely agreed to require time to propagate. Causes are generally assumed to precede their effects, and many social scientists, statisticians, and philosophers have claimed that time-ordering can be used for selecting appropriate causal models. However, ordinary practice is typically neither explicit nor careful about the relationship between causation and time. In our paper, we consider the role of temporal ordering in the practice of causal modeling, especially in the social sciences. We observe that in the social sciences, "time-free" causal models-i.e., ones which lack explicit information about the relative timing of causes and effects–are quite commonplace. One might hope that the time-free models proposed in practice will always (or for the most part) be consistent with the assumption that causes precede their effects. But we argue that things are not so tidy. Hence, we propose the following Temporalization Criterion of Admissibility for causal models: a causal model is *admissible* if and only if there is at least one temporalized version of the model—i.e., a version in which the variables are appropriately time-indexed—that both preserves a common-sense temporal ordering of causes and effects and also remains responsible to the data.

One might think that our criterion of admissibility could not possibly constrain actual modeling work. But we show that the criterion is non-trivial by pointing to examples of causal models in sociology (Simons et al. on religiosity and risky sexual behavior), political science (Timberlake and Williams on foreign investment and political exclusion), and education research (Beilock et al. on teacher math anxiety and student math achievement) that violate our criterion. Although our criterion applies to all causal models, we focus on time-free causal models, since for time-free models it is not always obvious that the criterion has been violated. A time-free causal model might have the appearance of reasonableness, but when one begins to think seriously and carefully about how to temporalize the model one encounters what we call the Temporalization Dilemma: one may select either a temporalized version of the model that has a common-sense temporal ordering of causes and effects or a temporalized version of the model that is responsible to the data, but not both. In other words, some time-free causal models appear reasonable but nonetheless violate the Temporalization Criterion of Admissibility. The upshot is that causal modelers need to think carefully about the temporal commitments of their models because in some cases, a model that looks reasonable despite lacking explicit temporal information cannot be temporalized in a way that both preserves a common-sense temporal ordering of causes and effects and also remains responsible to the data.

Thomas Lodewyckx and Bert Leuridan

Causation and Time: Synchronic Causality

It is the purpose of this presentation to present a case for causal relations of a synchronic nature. Causation is regularly and almost by definition assumed to be thoroughly diachronic or temporally ordered. Causes, as nearly any theory of causation will uphold, always precede their effects. However, it is our aim to investigate the possibility of scientifically acknowledged synchronic, or instantaneous, causal relations.

It is a generally held view by researchers in complex-system mechanisms, that causal relations are entirely restricted to intralevel relations. Interlevel relations, commonly typified by constitutive relevance or grounding relations, are then seen as devoid of causal efficacy (see Craver 2007 and Craver & Bechtel 2007). Building upon the suggestion made by Leuridan (2012), these interlevel relations could contrarily be regarded as a type of causal relations. Arguing on behalf of synchronicity could be seen as a first and necessary step in such a direction.

The argument we wish to present on synchronic causation is twofold. Firstly, we will analyze some examples of instantaneous causation at work in scientifically accepted processes. An instance of such an arguably synchronic causal process can be found in the reciprocal coupling of perception and action in cognitive agents, as described by Vernon et al. (2015). As evidenced by a number of neuroscientific and psychological experiments, perceptions influence actions and–vice versa–it can be said that actions change perceptions. In cognitive agents this relation is accompanied by certain constitutive qualities, such a self-organization and self-maintenance, giving us reason to believe that this relation may be constitutive in nature. If this is taken into account alongside the seemingly temporally simultaneous, bidirectionally causal relationship between perception and action, it could serve to support synchronic causation, as there is then both evidence of a constitutive and causal element in this relation. Causal relations in these examples will be interpreted using the groundwork provided by Woodward's interventionist theory of causality (2003/5), as his account plays an important role in the current debate surrounding constitutive relevance (Craver 2007). Secondly, we will focus on the more general argument that causality is supposedly thoroughly diachronic while constitutive relevance is contrastingly synchronic. To this end we examine some of the classic examples of synchronous causation provided by Heumer and Kovitz (2003). Though we will not subscribe to their eventual conclusion that all actual causes are simultaneous with their direct effects, there is some merit in recasting these examples in constitutive terms (though not necessarily from the perspective of Craver's MM-account of constitutive relevance) to evaluate the consequences for some of the typical critiques against the standard examples which they cite.

In the end, we intend to show that there are both good reasons to doubt

the entrenched aversion to synchronic causal relations and the supposed non-occurrence of synchronic causality in scientifically acknowledged processes. Besides making a case for the possibility of synchronic causality, this work might serve to further strengthen and elaborate on interventionist account of causation provided by Woodward (2003/5), by showing that it applies unproblematically to constitutive relations as well.

References

C. F. Craver (2007) Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience, Oxford: Oxford University Press.

C. F. Craver, W. Bechtel (2007) "Top-down Causation without Top-down Causes", Biology and Philosophy 22: 547-563.

B. Leuridan (2012) "Three Problems for the Mutual Manipulability Account of Constitutive Relevance in Mechanisms", The British Journal for the Philosophy of Science, 63(2): 399-427.

D. Vernon, R. Lowe, S. Thill and T. Ziemke (2015) "Embodied cognition and circular causality: on the role of constitutive autonomy in the reciprocal coupling of perception and action", Frontiers in Psychology 6: 1660.

J. Woodward (2003/5) Making Things Happen: A Theory of Causal Explanation, 2nd edition, Oxford: Oxford University Press.

M. Huemer, B. Kovitz (2003) "Causation as Simultaneous and Continuous", The Philosophical Quarterly 53: 213.

W. Han, L. Tellez, M. Rangel, et al. (2017) "Integrated Control of Predatory Hunting by the Central Nucleus of the Amygdala", Cell 168: 311-324.

Daniel Malinsky

Causal inference from time series data with unmeasured confounding: how data-driven can we get?

Algorithms which search for causal graphical models from observational data can be used to infer causal relations and inform decisions about interventions. While most applications of these methods have been in the i.i.d. domain, it is in many cases straightforward to adapt these algorithms to time series data. We discuss modeling dynamic systems with ancestral graph Markov models, which are well-suited to domains with possible unmeasured confounders (i.e., latent variables). Ancestral graphs can be thought of as marginalizations of DAG models; i.e., they represent causal relations and conditional independence facts even when some (unobserved) variables have been marginalized out. We give an overview of these ancestral graphical models for time series with an emphasis on their relationship to dynamic Bayesian networks and structural vector autoregressions (SVARs). We show how the usual methodology for learning SVARs in econometrics faces two limitations: the requirement of substantial background knowledge about causal relations to achieve identifiability, and the assumption of no unmeasured confounding. In contrast, graphical structure learning algorithms modified from the i.i.d. domain can overcome these limitations. We present constraint-based and score-based procedures for learning equivalence classes of (dynamic) ancestral graphs, and present an application to real economic data. We explore the prospects for relaxing some parametric assumptions like linearity, as well allowing for some kinds of heterogeneity. Though we can allow for non-linear relationships in the constraint-based approach, the score-based solution to the heterogeneity problem requires linearity, so there is currently some trade-off here. Finally, we discuss an open and important frontier in data-driven causal inference from time series: the problem of nonstationarity. We illustrate some of the pitfalls for causal inference from nonstationary data, and evaluate some possible solutions.

Flavia Padovani

At the Roots of the Common Cause

In his posthumous *The Direction of Time* (1956, sec. 19), Reichenbach formulated a principle that was bound to become one of the focal points of later discussions about causation: the principle of the common cause. According to Reichenbach, this principle "does not represent a new assumption, but is derivable from the second law of thermodynamics, if this law is supplemented by the hypothesis of the branch structure of thermodynamic systems." Nevertheless, this principle was already outlined in Reichenbach's early work, independently of those considerations. In "The Causal Structure of the World and the Difference between Past and Future" (1925), he introduces a rudimental version of the principle of the common cause as a development of his causal theory of time. The idea of that paper was to develop a topological account of the probabilistic implications that can be obtained from an analysis of the behaviour of interacting causal chains, in order to define the direction of time. It is in this paper that, for the very first time, Reichenbach puts forward a description of causal processes in terms of nets, and he presents the so-called fork asymmetry account on which he will indeed elaborate in (1956).

The model of topology that Reichenbach adopts in that early paper relies heavily on the temporal topology developed by a Gestalt psychologist, Kurt Lewin, in 1922-1923. In its basic form, in fact, Reichenbach's 1925 account builds quite consistently on Lewin's analysis of splitting and intersecting genetic series, that Lewin introduced primarily to illustrate the relationship between ancestors and progeny, but that he later expanded so as to include relations possessing a certain order (including the causal ones). One of these is called "genidentity", and it is defined as the relation between constructs that have derived one from the other. In Lewin's view, genidentity importantly gives rise to a specific type of order, namely the "existential relationship of the one-after-the-other"—an idea easily borrowed by Reichenbach. Following Lewin, Reichenbach emphasises that only the relations between actual events belonging to different (and ultimately, genidentical) series can provide good grounds for identifying the direction of causal chains. In this way, according to Reichenbach, "we are led to base the temporal order upon the characteristics of a net structure", as he will go on to do again in 1956.

The aim of the present paper is to examine Reichenbach's principle of the common cause in the light of his early work. I will illustrate how this principle is rooted in Reichenbach's causal theory of time and how it is deeply bound to the assumption of a principle of genidentity. I will show how some of the shortcomings related to this assumption in his early work are inherited in his later work, and from there to several current accounts that employ a derived form of the principle of the common cause.

Julian Reiss

Time Series, Non-Stationarity, and Causal Inference

Time series are data series indexed in time order where the ordering of the data points may matter.Most observational data are time series. A time series is stationary if and only if its moments — such as its mean and its variance — do not change over time. Thus, if the moments of a time series change over time, the time series is said to be non-stationary. Most time series in policy relevant sciences such as economics and climate science are non-stationary.

The difficulties for empirical modelling are vast when time series are non-stationary. If two causally unrelated time series are non-stationary because they evolve by accumulating past shocks, for instance, their correlation will appear to be significant at the 5% standard about 70% of the time. Examples include the correlation between murders and membership of the Church of England or that between the money stock and cumulative rainfall in the UK (the latter is due to econometrician David Hendry, 'Econometrics — Alchemy or Science?' *Economica* 1980). The physical and social processes underlying economic and climate phenomena all but ensures that most data in these fields are non-stationary.

Non-stationarity affects theorising, the estimation of parameters, forecasting, and policy. The purpose of this paper is to examine the implications of non-stationarity for causal inference. Everyone knows that 'correlation is not causation', but many methods of inference (for instance, Bayes nets, including dynamic Bayes nets, and regression analysis) assume that correlation is causation after all — after controlling for reverse causation (i.e., causation running from the putative effect to the putative cause) and 'third factors' or confounders. Non-stationarity shows that this conception of causation is untenable. Moreover, as non-stationarity is such a pervasive phenomenon, it shows that popular methods of inference are inapplicable except in very special circumstances.

On the positive side, I discuss a number of methods that work better in non-stationary contexts and, finally, draw some conclusions for our understanding of the concept of cause.

Naftali Weinberger

Can Causal Models Include Variables with Their Time-Derivatives?

Variables that are conceptually or mereologically related cause problems for causal inference. Learning that *being a bachelor* and *being unmarried* are "correlated" provides no evidence that one causes the other. Consequently, it is commonly assumed that the variables in a causal model are, in various ways, distinct from one another. Here I consider the relationship between a variable and its higher-order time-derivatives. These are not fully distinct. For instance, an object's present velocity has implications for its future position.

On the other hand, interventions on a variable are non-identical to those on its time-derivative; while ideal interventions on an object's position determine its position independently of its prior position, interventions on velocity determine the object's future trajectory in conjunction with its current position. We thus have reasons to consider models with variables and their time-derivatives. But we need to determine when modeling such conceptually-related variables is legitimate.

I begin with a debate about the causal role of instantaneous velocity in physics. Bertrand Russell reduces instantaneous velocity to the limit of an object's rate of change in the neighborhood surrounding the relevant instant. If so, an object's velocity at an instant is constituted by facts about its position at (arbitrarily close) moments both before and after that instant. Lange (2005) objects that if an object's velocity is a cause of all subsequent points in its trajectory, then we cannot define velocity in terms of its future trajectory without there being self-causation or backwards causation. One might try to evade Lange's problem by defining instantaneous velocity as its rate of change *prior* to that instant. This would preserve velocity's role as a cause, but at the expense of being able to treat it as an effect. Easwaran (2014) presents a way out of this dilemma, but at the expense of saying that instantaneous acceleration can only be an effect, never a cause. The lesson I draw from this debate is that one way to include both a variable and a derivative in a model is by stipulating that derivatives count only as effects, and that their influence on downstream variables is only via their being constituted by non-derivative variables.

After discussing these foundational issues, I consider "dynamic" causal models that use derivatives (Iwasaki and Simon, 1994; Dash, 2003). First, I note that these models satisfy the constraint that derivatives serve only as effects. This evades conceptual problems deriving from indistinctness. Second, I argue that such models need not be understood as literally representing instantaneous rates of change. Derivatives allow one to distinguish among variables that equilibrate at different rates in response to perturbations. Variables with higher-order derivatives equilibrate at slower rates than those represented without them. Under this interpretation, dynamic causal models are shorthand for discrete models with time-indexed variables that equilibrate at different rates.

References

Dash, Denver. 2003. "Caveats for Causal Reasoning with Equilibrium Models" (Thesis). University of Pittsburgh.

Easwaran, K. (2014). "Why physics uses second derivatives". *The British Journal for the Philosophy of Science*, 65(4), 845-862.

Iwasaki, Y., & Simon, H.A. (1994). "Causality and model abstraction". *Artificial Intelligence*, 67, 143-194. Lange, Marc. "How Can Instantaneous Velocity Fulfill Its Causal Role?" *Philosophical Review* (2005)

Michael Wilde

Another disambiguation of the Russo-Williamson thesis

Here is an epistemological thesis: At least in the health sciences, in order to establish a causal claim it is typically necessary to establish both that the putative cause and effect are appropriately correlated, and that they are linked by an appropriate mechanism. According to this thesis, in order to establish that aspirin causes relief from headaches, it is necessary to establish both that taking aspirin is correlated with headache relief, and that there exists a mechanism linking aspirin to headache relief that can explain the extent of this correlation. This thesis has been advocated by Federica Russo and Jon Williamson (2007; 2011). As a result, it is sometimes called the Russo-Williamson thesis (Gillies, 2011; Clarke et al., 2014).

The Russo-Williamson thesis has proved quite controversial (Weber, 2009; Broadbent, 2011; Campaner and Galavotti, 2012). However, it is likely that some of the controversy is the result of misinterpreting the thesis, since formulations of the thesis have been ambiguous in a number of different ways. This much has been argued by Phyllis Illari (2011). Illari proceeds to disambiguate the Russo-Williamson thesis. She argues that the thesis is quite plausible, as long as it is not interpreted as claiming that there is a distinctive method for establishing the existence of the underlying mechanism (2011:141–148). But even granting this point, it still looks like there are good objections to the Russo-Williamson thesis. In particular, Jeremy Howick (2011) has proposed a number of potential counterexamples, namely, some typical cases from the health sciences in which it looks like a causal claim was established without having established the existence of an underlying mechanism. For example, Howick claims that it was established that aspirin causes headache relief without having first established a mechanism linking aspirin to headache relief (2011:930).

In this paper, I provide a general response to these proposed counterexamples. I argue that they show only that establishing the existence of an underlying mechanism is not a necessary precondition for establishing the corresponding causal claim. However, this does not go against the Russo-Williamson thesis. The thesis is concerned with the necessary conditions rather than the necessary preconditions of establishing a causal claim. In other words, the thesis maintains that it is not possible to establish a causal claim without thereby establishing an appropriate mechanism. This response requires further disambiguating the Russo-Williamson thesis. And the disambiguation involves explicit reference to time. In particular, I argue that the Russo-Williamson thesis should not be interpreted as describing rational constraints that an agent's body of evidence must meet before a causal claim can be established. Instead, the thesis describes the rational constraints on an agent's body of evidence after they have established a causal claim, that is, what follows from an agent's having established a causal claim, namely, that they have thereby established

that there exists an underlying mechanism. Howick's proposed counterexamples miss their mark by failing to recognize this distinction.

References

Broadbent, A. (2011). Inferring causation in epidemiology: Mechanisms, black boxes, and contrasts. In P. M. Illari, F. Russo, and J. Williamson (Eds.), *Causality in the Sciences*, pp. 45–69. Oxford University Press.

Campaner, R. and M. C. Galavotti (2012). Evidence and the assessment of causal relations in the health sciences. *International Studies in the Philosophy of Science* 26(1), 27–45.

Clarke, B., D. Gillies, P. Illari, F. Russo, and J. Williamson (2014). Mechanisms and the evidence hierarchy. *Topoi* 33, 339–360.

Gillies, D. (2011). The Russo-Williamson thesis and the question of whether smoking causes heart disease. In P. M. Illari, F. Russo, and J. Williamson (Eds.), *Causality in the Sciences*, pp. 110–124. Oxford University Press.

Howick, J. (2011). Exposing the vanities-and a qualified defense-of mechanistic reasoning in health care decision making. *Philosophy of Science* 78, 926–940.

Illari, P. (2011). Mechanistic evidence: Disambiguating the Russo-Williamson thesis. *International Studies in the Philosophy of Science* 25, 139–157.

Russo, F. and J. Williamson (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science* 21, 157–170.

Russo, F. and J. Williamson (2011). Epistemic causality and evidence-based medicine. *History and Philosophy of the Life Sciences* 33, 563–582.

Weber, E. (2009). How probabilistic causation can account for the use of mechanistic evidence. *International Studies in the Philosophy of Science* 23, 277–295.

Yi Zhao and Xi Luo

Granger Mediation Analysis of Functional Magnetic Resonance Imaging Time Series

Making inference about brain effective connectivity is of great interest in task functional magnetic resonance imaging (fMRI) experiments. In this study, we are interested in clarifying the causal mechanisms of an external stimulus on the interested outcome brain region, by considering another brain region that functions as an intermediate variable. To achieve this, causal mediation analysis under structural equation modeling framework is considered. However, attaining causal interpretations requires both the "no unmeasured confounding" and the "no interference" assumptions. These two assumptions generally do not hold in fMRI datasets. To address the existence of unmeasured confounding, a correlation between model errors is introduced; and to characterize the temporal and interregional dependency, the principle of Granger causality is implemented. In this paper, we propose a Granger Mediation Analysis framework that provides inference about both spatial and temporal causality between brain regions for multilevel fMRI time series. Simulation studies show that our method reduces the bias in estimating the causal effects compared to existing approaches. Applying the proposed method on a real fMRI dataset, our approach not only estimates the causal effects of brain pathways, but effectively captures the feedback effect of the outcome region on the mediator region.

About Stevens Institute of Technology

Stevens Institute of Technology, The Innovation University, is a premier, private research university situated in Hoboken, N.J. overlooking the Manhattan skyline. Founded in 1870, technological innovation has been the hallmark and legacy of Stevens' education and research programs for 147 years. Within the university's three schools and one college, 6,600 undergraduate and graduate students collaborate with more than 300 full-time faculty members in an interdisciplinary, student-centric, entrepreneurial environment to advance the frontiers of science and leverage technology to confront global challenges. Stevens is home to three national research centers of excellence, as well as joint research programs focused on critical industries such as healthcare, energy, finance, defense, maritime security, STEM education and coastal sustainability. For more information, please visit stevens.edu