

Intervention & Observation

ABNMS 2011

Intervention & Observation

Kevin B. Korb
Clayton School of Information Technology
Monash University
kbborb@gmail.com

Intervention & Observation

Abstract

Bayesian nets were originally designed to represent the effects of observation under uncertainty. They have sometimes been used to represent the effects of intervention – often misleadingly, since interventions and observations have different probabilistic consequences. I will review the differences and how Bayesian nets (and decision nets) can, and cannot, be used to represent causal interventions.

Fisher's Smoking Model

The difference between intervention and observation is clear.

E.g., as an alternative to *Smoking* → *Cancer*, Fisher (1957) proposed a common genetic factor.

Fisher's Smoking Model

The difference between intervention and observation is clear.

E.g., as an alternative to *Smoking* \rightarrow *Cancer*, Fisher (1957) proposed a common genetic factor.

$$\textit{Smoking} \leftarrow \textit{Gene} \rightarrow \textit{Cancer}$$

Fisher's Smoking Model

The difference between intervention and observation is clear.

E.g., as an alternative to *Smoking* \rightarrow *Cancer*, Fisher (1957) proposed a common genetic factor.

$$\textit{Smoking} \leftarrow \textit{Gene} \rightarrow \textit{Cancer}$$

- ▶ Observation of *Smoking* affects *Cancer*

Fisher's Smoking Model

The difference between intervention and observation is clear.

E.g., as an alternative to *Smoking* \rightarrow *Cancer*, Fisher (1957) proposed a common genetic factor.

$$\textit{Smoking} \leftarrow \textit{Gene} \rightarrow \textit{Cancer}$$

- ▶ Observation of *Smoking* affects *Cancer*
- ▶ Intervention on *Smoking* does not affect *Cancer*
(that was Fisher's point!)

Fisher's Smoking Model

Was Fisher right that his model endorses smoking?

Fisher's Smoking Model

Was Fisher right that his model endorses smoking?

Smoking ← *Gene* → *Cancer*

Fisher's Smoking Model

Was Fisher right that his model endorses smoking?

$$\textit{Smoking} \leftarrow \textit{Gene} \rightarrow \textit{Cancer}$$

Supposing: (a) Fisher's model is right; and (b) you

- ▶ Really hate getting cancer
- ▶ But mildly like smoking, other things being equal

Fisher's Smoking Model

Was Fisher right that his model endorses smoking?

$$\textit{Smoking} \leftarrow \textit{Gene} \rightarrow \textit{Cancer}$$

Supposing: (a) Fisher's model is right; and (b) you

- ▶ Really hate getting cancer
- ▶ But mildly like smoking, other things being equal

Should you smoke?

Fisher's Smoking Model

Principle (Maximize Expected Utility (MEU))

Given k options, you should take one that maximizes your expected utility.

Fisher's Smoking Model

Principle (Maximize Expected Utility (MEU))

Given k options, you should take one that maximizes your expected utility.

Smoking \leftarrow *Gene* \rightarrow *Cancer*

Fisher's Smoking Model

Principle (Maximize Expected Utility (MEU))

Given k options, you should take one that maximizes your expected utility.

Smoking \leftarrow Gene \rightarrow Cancer

The expected utility of smoking v not smoking:

$$P(C|S)u(C|S) + P(\neg C|S)u(\neg C|S) \ll P(C|\neg S)u(C|\neg S) + P(\neg C|\neg S)u(\neg C|\neg S)$$

Fisher's Smoking Model

Principle (Maximize Expected Utility (MEU))

Given k options, you should take one that maximizes your expected utility.

Smoking \leftarrow Gene \rightarrow Cancer

The expected utility of smoking v not smoking:

$$P(C|S)u(C|S) + P(\neg C|S)u(\neg C|S) \ll P(C|\neg S)u(C|\neg S) + P(\neg C|\neg S)u(\neg C|\neg S)$$

\Rightarrow *Evidential decision theory: Don't smoke!*

Fisher's Smoking Model

Principle (Pareto Dominance)

Given k options, if one of them has outcomes everywhere at least as good as all other outcomes and somewhere better, then you should take that option.

Fisher's Smoking Model

Principle (Pareto Dominance)

Given k options, if one of them has outcomes everywhere at least as good as all other outcomes and somewhere better, then you should take that option.

Smoking \leftarrow *Gene* \rightarrow *Cancer*

Fisher's Smoking Model

Principle (Pareto Dominance)

Given k options, if one of them has outcomes everywhere at least as good as all other outcomes and somewhere better, then you should take that option.

Smoking \leftarrow *Gene* \rightarrow *Cancer*

Smoking dominates. Either you already have G or you don't, so:

- ▶ G + Smoking:
$$P(C|G)u(C, S) + P(\neg C|G)u(\neg C, S) > P(C|G)u(C, \neg S) + P(\neg C|G)u(\neg C, \neg S)$$
- ▶ $\neg G$ + Smoking: $P(C|\neg G)u(C, S) + P(\neg C|\neg G)u(\neg C, S) >$
$$P(C|\neg G)u(C, \neg S) + P(\neg C|\neg G)u(\neg C, \neg S)$$

Fisher's Smoking Model

Principle (Pareto Dominance)

Given k options, if one of them has outcomes everywhere at least as good as all other outcomes and somewhere better, then you should take that option.

Smoking \leftarrow *Gene* \rightarrow *Cancer*

Smoking dominates. Either you already have G or you don't, so:

- ▶ G + Smoking:
$$P(C|G)u(C, S) + P(\neg C|G)u(\neg C, S) > P(C|G)u(C, \neg S) + P(\neg C|G)u(\neg C, \neg S)$$
- ▶ $\neg G$ + Smoking: $P(C|\neg G)u(C, S) + P(\neg C|\neg G)u(\neg C, S) >$
$$P(C|\neg G)u(C, \neg S) + P(\neg C|\neg G)u(\neg C, \neg S)$$

\Rightarrow *Causal decision theory: Go ahead and smoke!*

Decision/Intervention v Observation

So,

Decision/Intervention v Observation

So,

- ▶ Are actions, choices, decisions, interventions matters of *free will*, utterly and completely independent of their causal context and so properly modeled as BN Decisions and/or Pearl's do-calculus (Pearl, 2000)?

Decision/Intervention v Observation

So,

- ▶ Are actions, choices, decisions, interventions matters of *free will*, utterly and completely independent of their causal context and so properly modeled as BN Decisions and/or Pearl's do-calculus (Pearl, 2000)?
- ▶ Or, are humans and human behavior no different causally from any other natural process and so properly modeled as BN Chance nodes?

Decision/Intervention v Observation

So,

- ▶ Are actions, choices, decisions, interventions matters of *free will*, utterly and completely independent of their causal context and so properly modeled as BN Decisions and/or Pearl's do-calculus (Pearl, 2000)?
- ▶ Or, are humans and human behavior no different causally from any other natural process and so properly modeled as BN Chance nodes?

Decision/Intervention v Observation

So,

- ▶ Are actions, choices, decisions, interventions matters of *free will*, utterly and completely independent of their causal context and so properly modeled as BN Decisions and/or Pearl's do-calculus (Pearl, 2000)?
- ▶ Or, are humans and human behavior no different causally from any other natural process and so properly modeled as BN Chance nodes?

Are humans angels or animals?

Newcomb's Paradox

Nozick (1969)

An eccentric millionaire psychologist proposes a game in which you could win \$1M:

Newcomb's Paradox

Nozick (1969)

An eccentric millionaire psychologist proposes a game in which you could win \$1M:

- ▶ She has studied your psychology and can make perfect predictions of your future behavior (and she's right & you believe her)

Newcomb's Paradox

Nozick (1969)

An eccentric millionaire psychologist proposes a game in which you could win \$1M:

- ▶ She has studied your psychology and can make perfect predictions of your future behavior (and she's right & you believe her)
- ▶ She puts two boxes in front of you. You can choose one box (the right box) or both.
 - ▶ Left box: \$1
 - ▶ Right box: \$1M if you are predicted to choose one box; \$0 if you are predicted to choose two boxes.

Newcomb's Paradox

Nozick (1969)

An eccentric millionaire psychologist proposes a game in which you could win \$1M:

- ▶ She has studied your psychology and can make perfect predictions of your future behavior (and she's right & you believe her)
- ▶ She puts two boxes in front of you. You can choose one box (the right box) or both.
 - ▶ Left box: \$1
 - ▶ Right box: \$1M if you are predicted to choose one box; \$0 if you are predicted to choose two boxes.

Newcomb's Paradox

Nozick (1969)

An eccentric millionaire psychologist proposes a game in which you could win \$1M:

- ▶ She has studied your psychology and can make perfect predictions of your future behavior (and she's right & you believe her)
- ▶ She puts two boxes in front of you. You can choose one box (the right box) or both.
 - ▶ Left box: \$1
 - ▶ Right box: \$1M if you are predicted to choose one box; \$0 if you are predicted to choose two boxes.

Should you one box or two box?

Newcomb's Paradox

Causal Decision Theory

The correct model is:

Choose Boxes \leftarrow *Psychology* \rightarrow *Psychologist* \rightarrow \$

Newcomb's Paradox

Causal Decision Theory

The correct model is:

Choose Boxes \leftarrow *Psychology* \rightarrow *Psychologist* \rightarrow \$

- ▶ The choice, whether One or Two Boxing, should be modeled as a Decision/Intervention/Free Will/Angelic Independence, cutting the link to your Psychology.

Newcomb's Paradox

Causal Decision Theory

The correct model is:

Choose Boxes \leftarrow *Psychology* \rightarrow *Psychologist* \rightarrow \$

- ▶ The choice, whether One or Two Boxing, should be modeled as a Decision/Intervention/Free Will/Angelic Independence, cutting the link to your Psychology.
- ▶ So, Two Boxing dominates.

Newcomb's Paradox

Evidential Decision Theory

The correct model is:

Choose Boxes \leftarrow *Psychology* \rightarrow *Psychologist* \rightarrow \$

Newcomb's Paradox

Evidential Decision Theory

The correct model is:

Choose Boxes \leftarrow *Psychology* \rightarrow *Psychologist* \rightarrow \$

- ▶ Two Boxing indicates the kind of psychology that yields poor results.

Newcomb's Paradox

Evidential Decision Theory

The correct model is:

Choose Boxes \leftarrow *Psychology* \rightarrow *Psychologist* \rightarrow \$

- ▶ Two Boxing indicates the kind of psychology that yields poor results.
- ▶ MEU applies, so you should One Box.

Evidential v Causal Decision Theory

Meeks & Glymour (1994): The disagreement is *not* between the two principles — Dominance and MEU. The disagreement is just about what the right model of decision making is.


Evidential v Causal Decision Theory

Meeks & Glymour (1994): The disagreement is *not* between the two principles — Dominance and MEU. The disagreement is just about what the right model of decision making is.

EU and Dominance do not, and can not, disagree!


Fisher's Smoking Model

Evidential Decision Theory



Fisher's Smoking Model


Evidential Decision Theory



- ▶ An *observation* of smoking raises the probability of cancer, lowering EU.


Fisher's Smoking Model

Evidential Decision Theory



Fisher's Smoking Model


Evidential Decision Theory



- ▶ An *observation* of not smoking lowers the probability of cancer, raising EU.


Fisher's Smoking Model

Perfect Intervention (Pearl, 2000)



Fisher's Smoking Model


Perfect Intervention (Pearl, 2000)



- ▶ A *perfect (Pearlian) intervention* to force smoking cuts the link to G, renders *Cancer* independent of smoking, raising EU.


Fisher's Smoking Model

Perfect Intervention



Fisher's Smoking Model


Perfect Intervention



- ▶ A *perfect (Pearlian) intervention* to force not smoking cuts the link to G, renders *Cancer* independent of smoking, lowering EU.


Fisher's Smoking Model

Imperfect Intervention



Fisher's Smoking Model


Imperfect Intervention



- ▶ An *imperfect intervention* for smoking does *not* cut the link to G. *ISmoke* and *Cancer* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.

Fisher's Smoking Model


Imperfect Intervention



- ▶ An *imperfect intervention* for smoking does *not* cut the link to G. *ISmoke* and *Cancer* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.
- ▶ No simple answers.


Fisher's Smoking Model

Imperfect Intervention



Fisher's Smoking Model


Imperfect Intervention



- ▶ An *imperfect intervention* for not smoking does *not* cut the link to G. *ISmoke* and *Cancer* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.

Fisher's Smoking Model

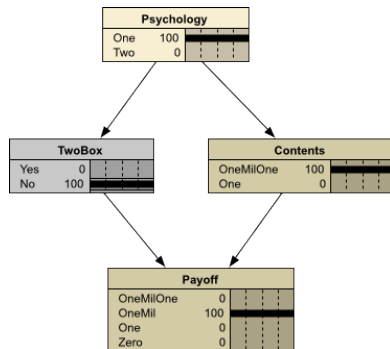
Imperfect Intervention



- ▶ An *imperfect intervention* for not smoking does *not* cut the link to G. *ISmoke* and *Cancer* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.
- ▶ No simple answers.

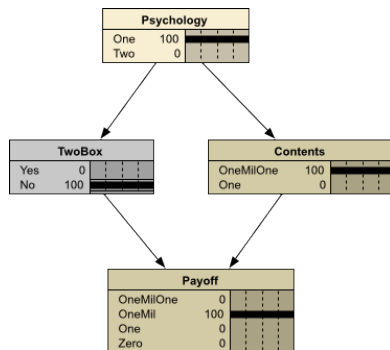
Newcomb's Paradox

Evidential Decision Theory



Newcomb's Paradox

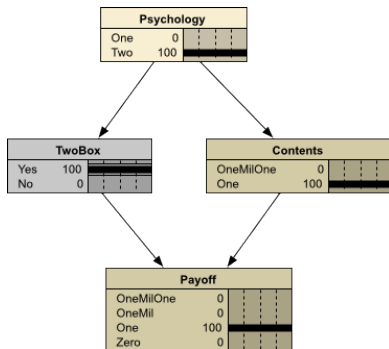
Evidential Decision Theory



- ▶ An *observation* of One Boxing raises the probability of \$1M, raising EU.

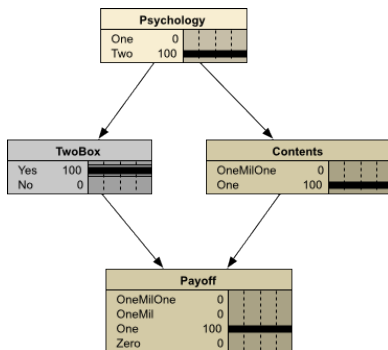
Newcomb's Paradox

Evidential Decision Theory



Newcomb's Paradox


Evidential Decision Theory



- ▶ An *observation* of Two Boxing lowers the probability of \$1M1, lowering EU.

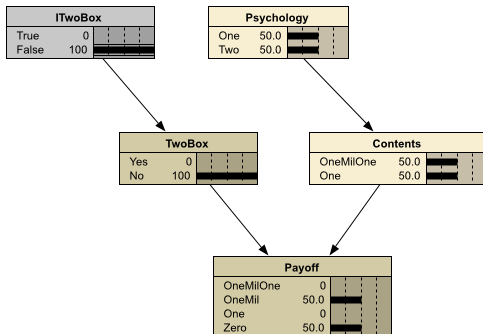
Newcomb's Paradox

Causal Decision Theory



Newcomb's Paradox

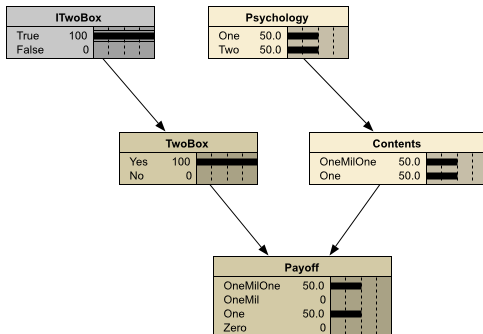
Causal Decision Theory



- ▶ A *perfect (Pearlian) intervention* to One Box cuts the link to *Psychology*, renders *Contents* independent of *TwoBox*, lowering EU.


Newcomb's Paradox

Causal Decision Theory



Newcomb's Paradox

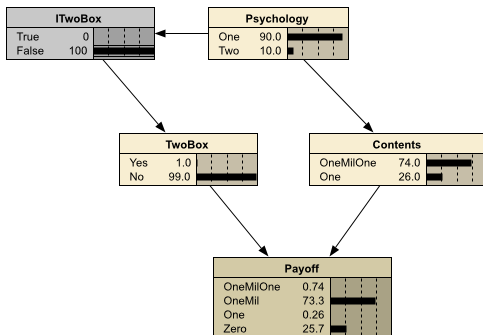
Causal Decision Theory



- ▶ A perfect (Pearlian) intervention to Two Box cuts the link to Psychology, renders Contents independent of TwoBox, raising EU.

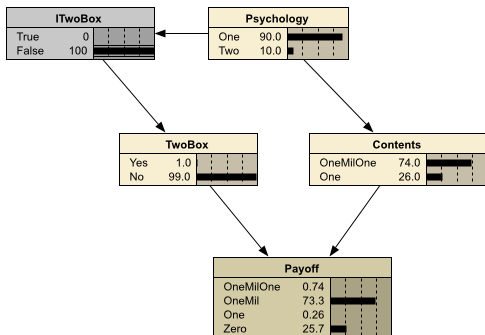
Newcomb's Paradox

Imperfect Intervention



Newcomb's Paradox

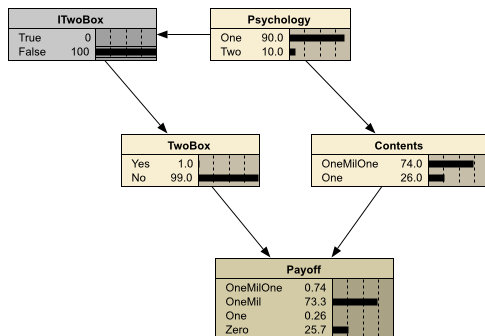
Imperfect Intervention



- ▶ An *imperfect intervention* to One Box does *not* cut the link to *Psychology*. *ITwoBox* and *Contents* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.

Newcomb's Paradox

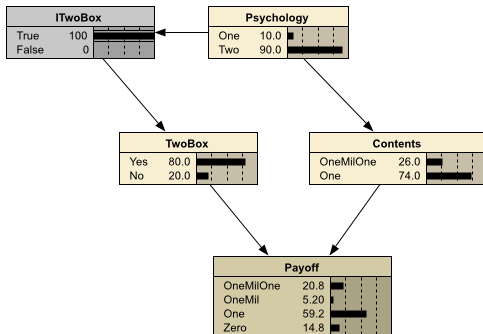
Imperfect Intervention



- ▶ An *imperfect intervention* to One Box does *not* cut the link to *Psychology*. *ITwoBox* and *Contents* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.
- ▶ No simple answers.

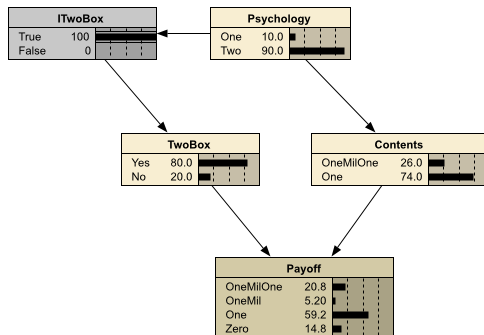
Newcomb's Paradox

Imperfect Intervention



Newcomb's Paradox

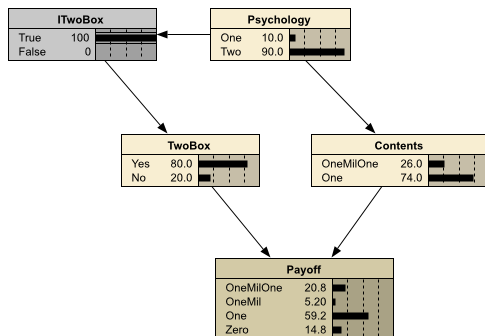
Imperfect Intervention



- ▶ An *imperfect intervention* to Two Box does *not* cut the link to *Psychology*. *ITwoBox* and *Contents* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.

Newcomb's Paradox

Imperfect Intervention



- ▶ An *imperfect intervention* to Two Box does *not* cut the link to *Psychology*. *ITwoBox* and *Contents* are dependent (Korb, et al., 2004). EU depends upon the *degree* of dependence.
- ▶ No simple answers.

Arguments

- ▶ CDT: A decision cannot affect the past, so Pearl's do-calculus should be used to compute EU.

Arguments

- ▶ CDT: A decision cannot affect the past, so Pearl's do-calculus should be used to compute EU.
- ▶ EDT: A decision cannot affect the past, but it can *reflect* the past. To the extent that it does so, imperfect intervention models should be used to compute EU.

Arguments

- ▶ CDT: A decision cannot affect the past, so Pearl's do-calculus should be used to compute EU.
- ▶ EDT: A decision cannot affect the past, but it can *reflect* the past. To the extent that it does so, imperfect intervention models should be used to compute EU.
- ▶ CDT: A decision may reflect the past, but *rational* decision making should do so as little as possible. The past is old baggage.

Arguments

- ▶ CDT: A decision cannot affect the past, so Pearl's do-calculus should be used to compute EU.
- ▶ EDT: A decision cannot affect the past, but it can *reflect* the past. To the extent that it does so, imperfect intervention models should be used to compute EU.
- ▶ CDT: A decision may reflect the past, but *rational* decision making should do so as little as possible. The past is old baggage.
 - ▶ For example, throwing good money after bad is a classic fallacy, failing to maximize EU.

Arguments

- ▶ CDT: A decision cannot affect the past, so Pearl's do-calculus should be used to compute EU.
- ▶ EDT: A decision cannot affect the past, but it can *reflect* the past. To the extent that it does so, imperfect intervention models should be used to compute EU.
- ▶ CDT: A decision may reflect the past, but *rational* decision making should do so as little as possible. The past is old baggage.
 - ▶ For example, throwing good money after bad is a classic fallacy, failing to maximize EU.

Arguments

- ▶ CDT: A decision cannot affect the past, so Pearl's do-calculus should be used to compute EU.
- ▶ EDT: A decision cannot affect the past, but it can *reflect* the past. To the extent that it does so, imperfect intervention models should be used to compute EU.
- ▶ CDT: A decision may reflect the past, but *rational* decision making should do so as little as possible. The past is old baggage.
 - ▶ For example, throwing good money after bad is a classic fallacy, failing to maximize EU.

We may *wish* we had a One-Boxer Psychology and collect \$1M; we should *still* Two Box.

Arguments

- ▶ CDT: A decision cannot affect the past, so Pearl's do-calculus should be used to compute EU.
- ▶ EDT: A decision cannot affect the past, but it can *reflect* the past. To the extent that it does so, imperfect intervention models should be used to compute EU.
- ▶ CDT: A decision may reflect the past, but *rational* decision making should do so as little as possible. The past is old baggage.
 - ▶ For example, throwing good money after bad is a classic fallacy, failing to maximize EU.

We may *wish* we had a One-Boxer Psychology and collect \$1M; we should *still* Two Box.

- ▶ Realistic Decision Theory: We are neither Angels nor Nietzschean Superhumans. To the extent we can figure out what we are, we should model ourselves correctly and use those models to maximize EU.

Conclusions

- ▶ Observations v Decisions/Interventions are radically different.

Conclusions

- ▶ Observations v Decisions/Interventions are radically different.
- ▶ Pearl's do-calculus (perfect interventions) can be a useful simplification; they usually misrepresent a more complex reality, however. We shouldn't expect ourselves to be *able* to instantiate the Pearlian ideal. (We are not Angels.)

Conclusions

- ▶ Observations v Decisions/Interventions are radically different.
- ▶ Pearl's do-calculus (perfect interventions) can be a useful simplification; they usually misrepresent a more complex reality, however. We shouldn't expect ourselves to be *able* to instantiate the Pearlian ideal. (We are not Angels.)
- ▶ Modeling imperfect interventions can be done by adding ordinary Chance nodes and determining their interdependencies with the rest of the causal nexus.

Conclusions

- ▶ Observations v Decisions/Interventions are radically different.
- ▶ Pearl's do-calculus (perfect interventions) can be a useful simplification; they usually misrepresent a more complex reality, however. We shouldn't expect ourselves to be *able* to instantiate the Pearlian ideal. (We are not Angels.)
- ▶ Modeling imperfect interventions can be done by adding ordinary Chance nodes and determining their interdependencies with the rest of the causal nexus.
- ▶ MEU and Dominance are not in conflict. All the conflicts are about finding the right causal model.

FIN

References

- Fisher, R.A. (1957). Dangers of cigarette-smoking. Letter in the *British Medical Journal*, 2, 297-298.
- Korb, K.B., Hope, L., Nicholson, A. and Axnick, K. (2004) Varieties of causal intervention. *Pacific Rim Conference on Artificial Intelligence: Trends in Artificial Intelligence*, 322–331. Springer Verlag.
- Meek, C. and Glymour, C. (1994). Conditioning and intervening. *The British Journal for the Philosophy of Science*, 45, 1001–1021.
- Nozick, R. (1969). Newcomb's Problem and Two Principles of Choice. In N. Rescher et al. (Eds.) *Essays in Honor of Carl C. Hempel*, pp. 114–146. Dordrecht Reidel.
- Pearl, J. (2000). *Causality: Models, Reasoning and Inference*. Cambridge Univ Press.