AI & Law models of case-based reasoning and their applications

Henry Prakken

IC3K'25, Marbella

22 October 2025

Utrecht University

ICS Department of Information & Computing Sciences

Marek Sergot

Institutional history

Thorne McCarty

- Florence conferences 1982,1985,1989
- ICAIL conferences since 1987
- JURIX conferences since 1988
- AI & Law journal since 1992
- ...
- Two landmark papers:
  - Taxman (Thorne McCarty, 1977): precedents
  - British Nationality Act (Marek Sergot et al., 1985): legislation

Factor-based reasoning

- In legal classification and interpretation there are often no clear rules
- Often there only are factors: tentative reasons pro or con a conclusion
  - Often to different degrees
- Factors are weighed in cases, which become precedents
  - How do judges, sollicitors … then argue?
  - And how do precedents constrain new decisions?

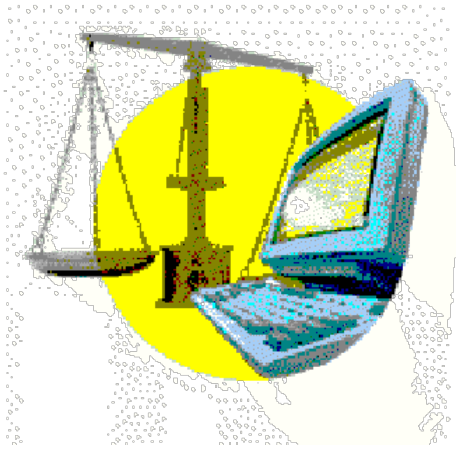

Example from US law: misuse of trade secrets

- *Some factors pro misuse of trade secrets*:
  - F4 Agreed-Not-To-Disclose
  - F6 Security-Measures
  - F18 Identical-Products
  - F21 Knew-Info-Confidential
  - …
- *Some factors con misuse of trade secrets*:
  - F1 Disclosure-In-Negotiations
  - F16 Info-Reverse-Engineerable
  - …

HYPO
Ashley & Rissland
1985-1990

CATO
Aleven & Ashley
1991-1997

HYPO

Ashley & Rissland 1987-1990

Kevin Ashley &
Edwina Rissland

- Representation language:
  - Cases: π-factors and δ-factors + decision (π or δ)
  - Current Fact Situation: π-factors and δ-factors
- Arguments:
  - Citing (for its decision) a case on its similarities with CFS
  - Distinguishing a case on its differences with CFS
  - Taking into account which side is favoured by a factor

Example with factors

C_1 (π)

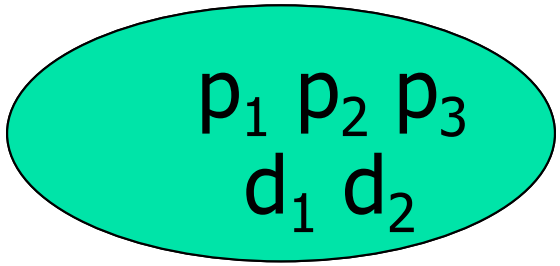

C_2 (δ)

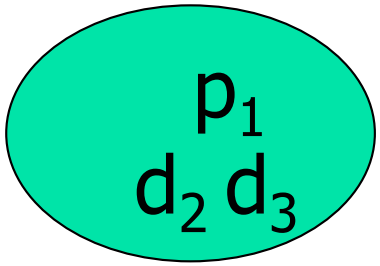

CFS

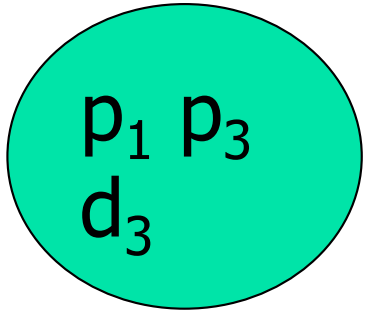

Example with factors

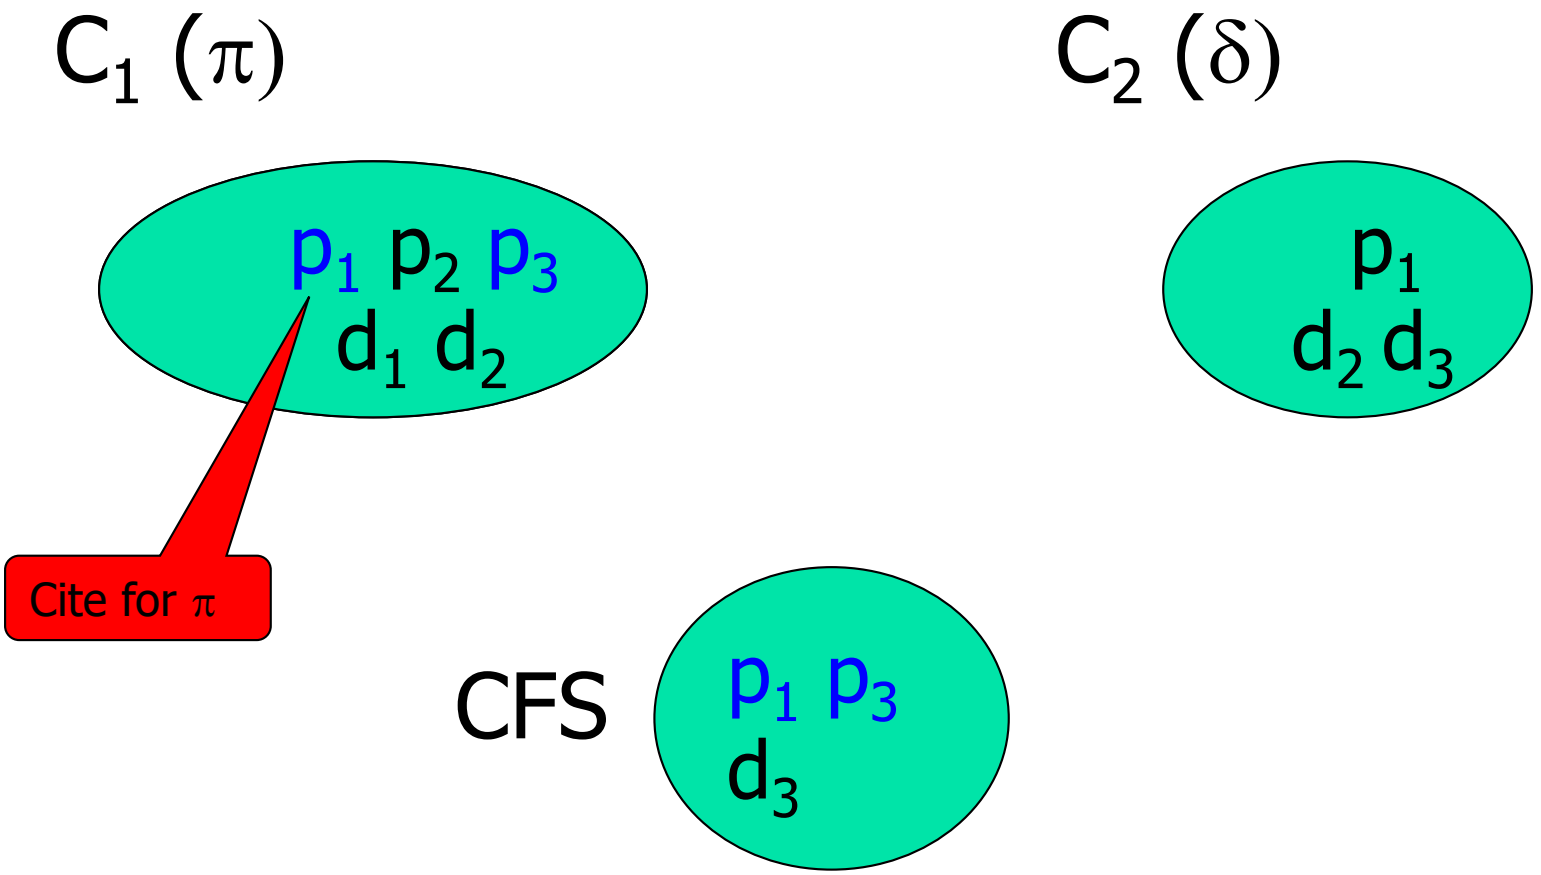

Example with factors

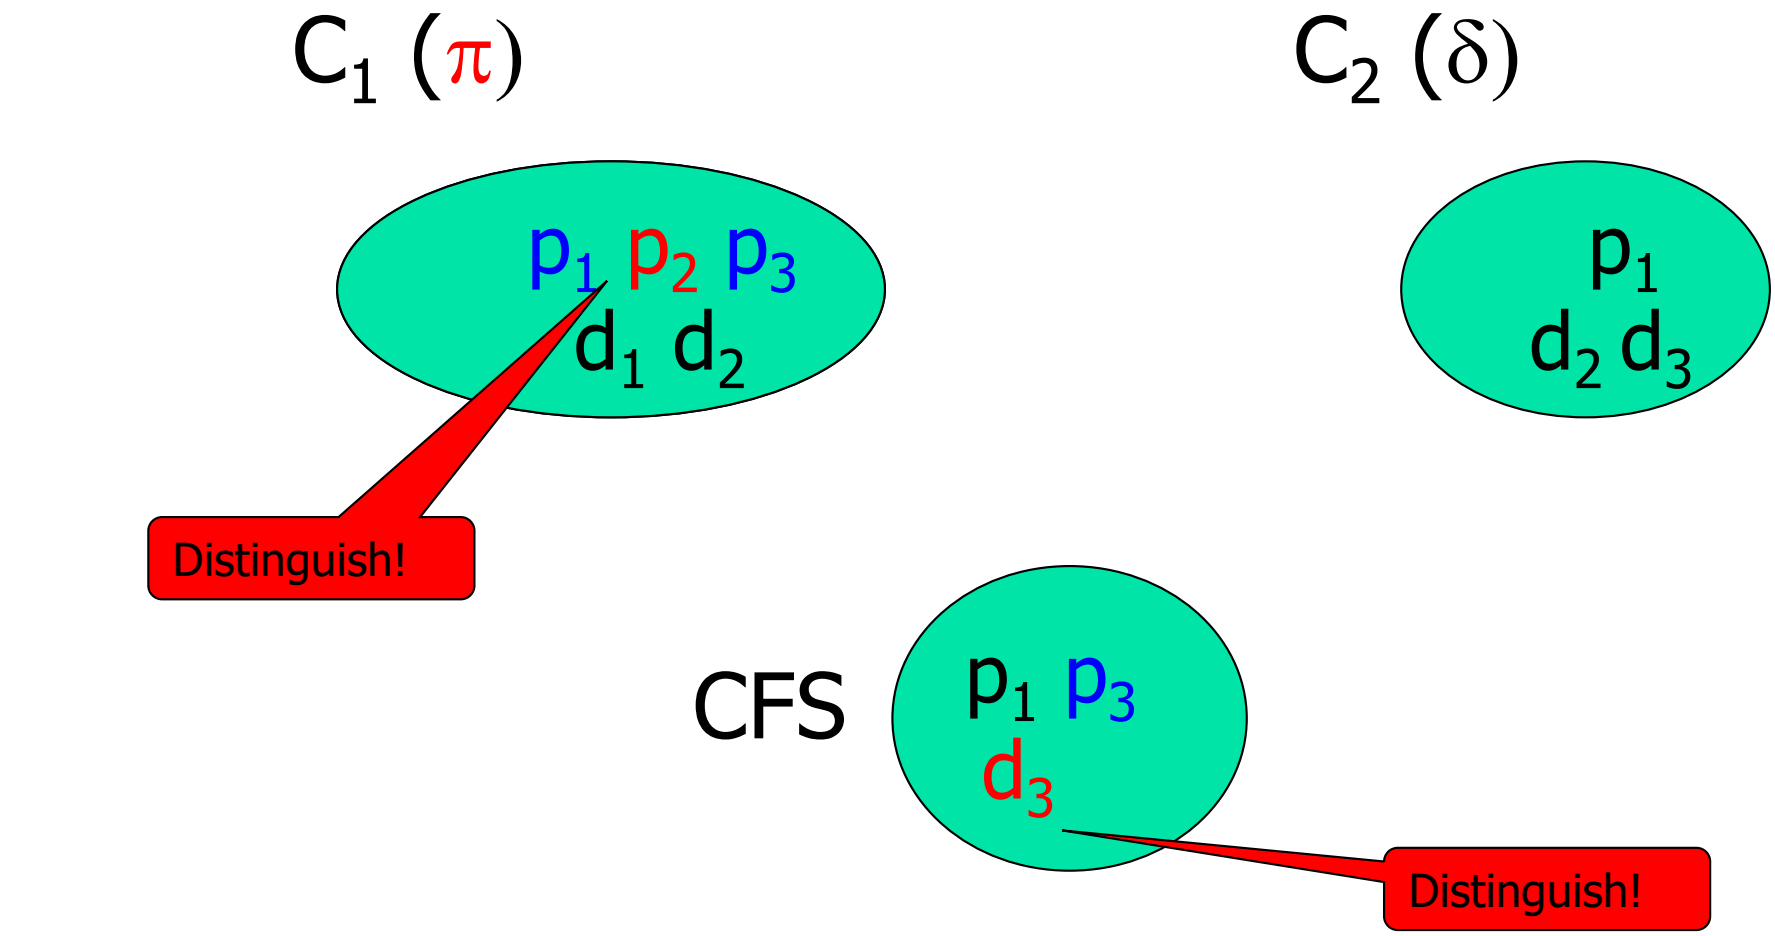

Example with factors

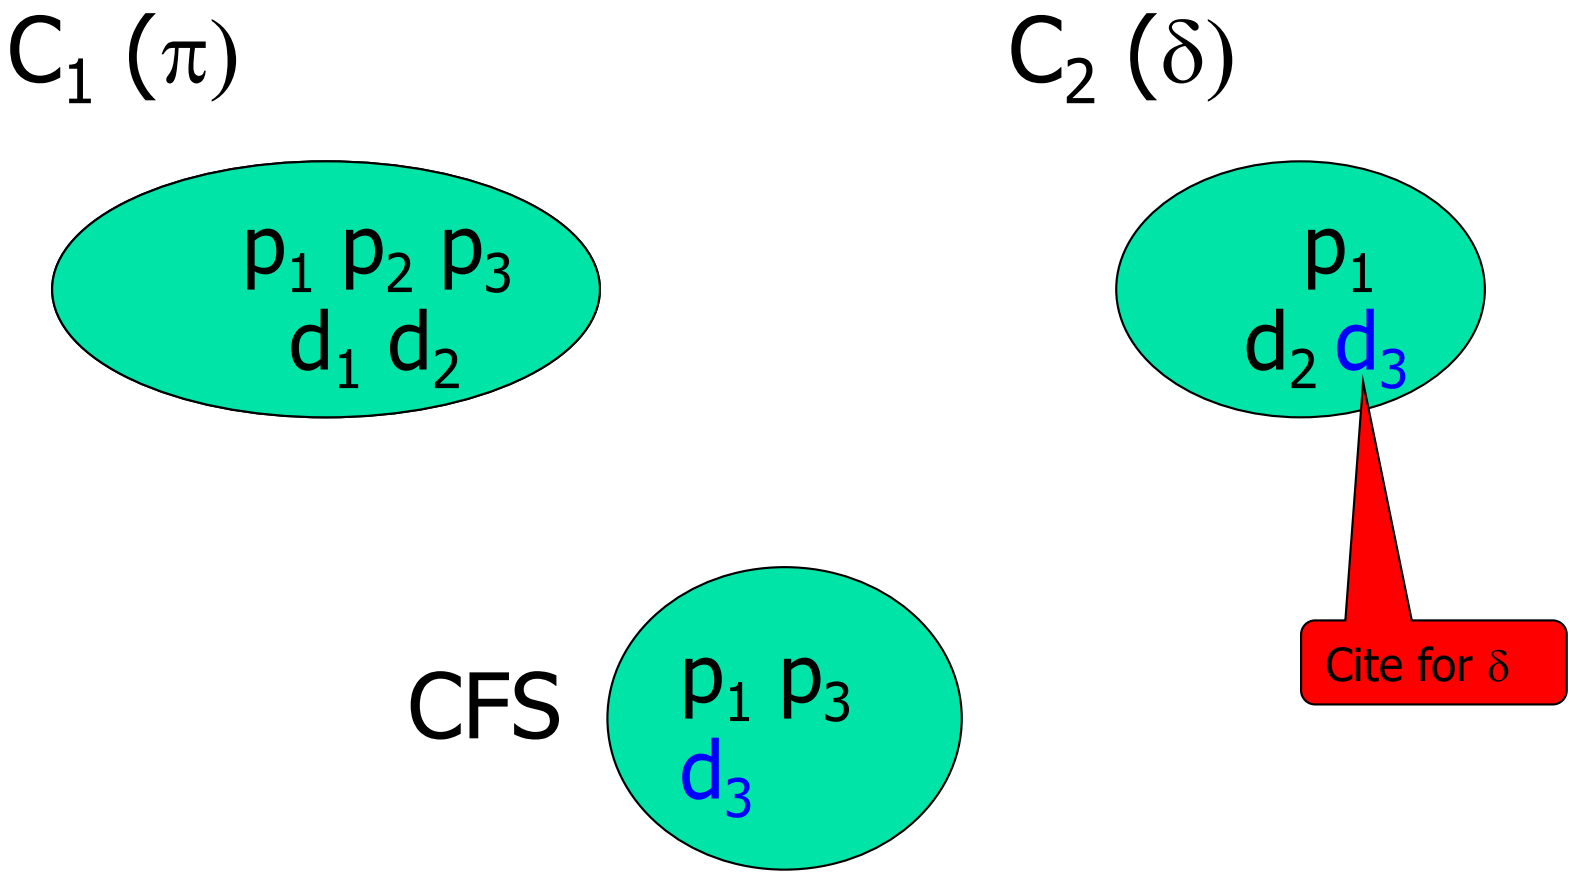

Example with factors

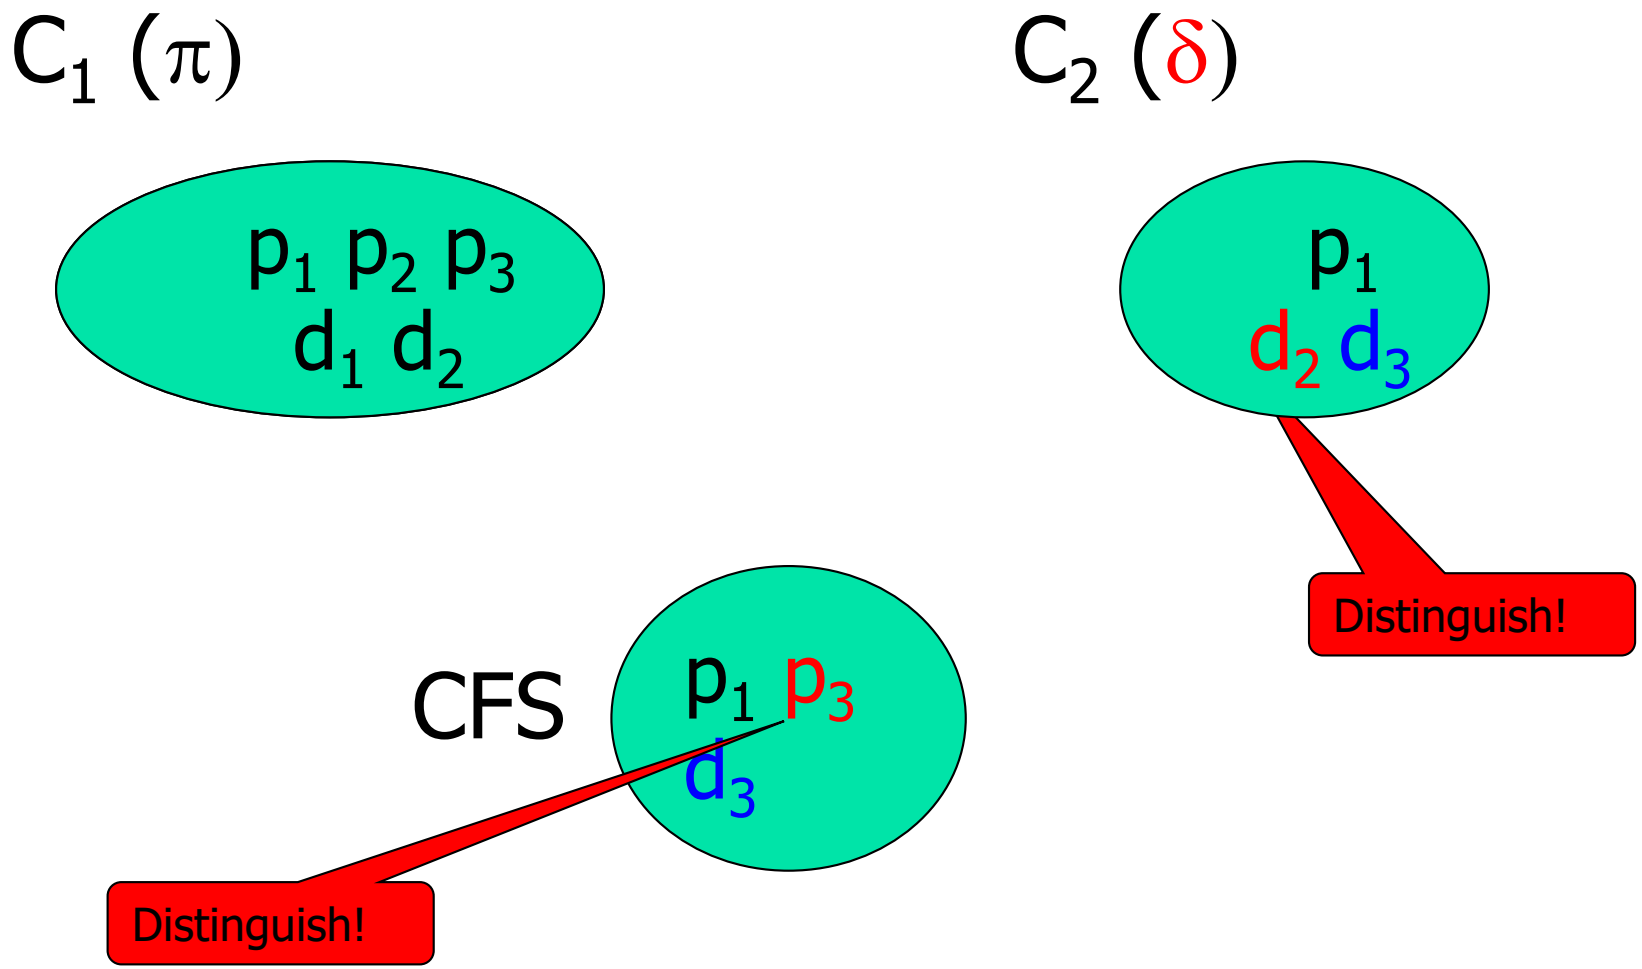

HYPO's argument game

- Given: a case base and a current fact situation
- Plaintiff starts with a citable case
  - A case decided for plaintiff and sharing pro-plaintiff factors with the CFS
- Defendant:
  - cites all counterexamples (cases citable for defendant)
  - distinguishes citation in all possible ways
    - On pro-plaintiff factors of precedent lacking in CFS
    - On new pro-defendant factors in the CFS
- Plaintiff distinguishes defendant's counterexamples in all possible ways

K.D. Ashley. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. MIT Press, Cambridge, MA, 1990.

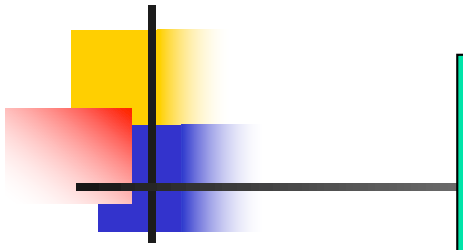

Plaintiff:
I should win because as in Bryce, which was won by plaintiff, I took security measures and defendant knew the info was confidential

Defendant:
Unlike in the present case, in Bryce defendant had agreed not to disclose and the products were identical

Defendant:
Unlike Bryce, in the present case the info is reverse engineerable

Defendant:
I should win because as in Robinson, which was won by defendant, plaintiff made disclosures during the negotiations

Plaintiff:
Unlike in Robinson, I took security measures, and defendant knew the info was confidential

Vincent Aleven 1991-1997

(snapshot of) CATO Factor Hierarchy

Misuse of Trade Secret (p)

F120: Info legitimately obtained elsewhere (d)

F101: Info Trade Secret (p)

F102: Efforts to maintain secrecy (p)

F104: Info valuable (p)

F4: Agreed not to disclose (p)

F1: Disclosures in negotiations (d)

F6: Security measures (p)

F15: Unique product (p)

V. Aleven. Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment. *Artificial Intelligence* 150:183-237, 2003.

Distinguishing on missing pro factor

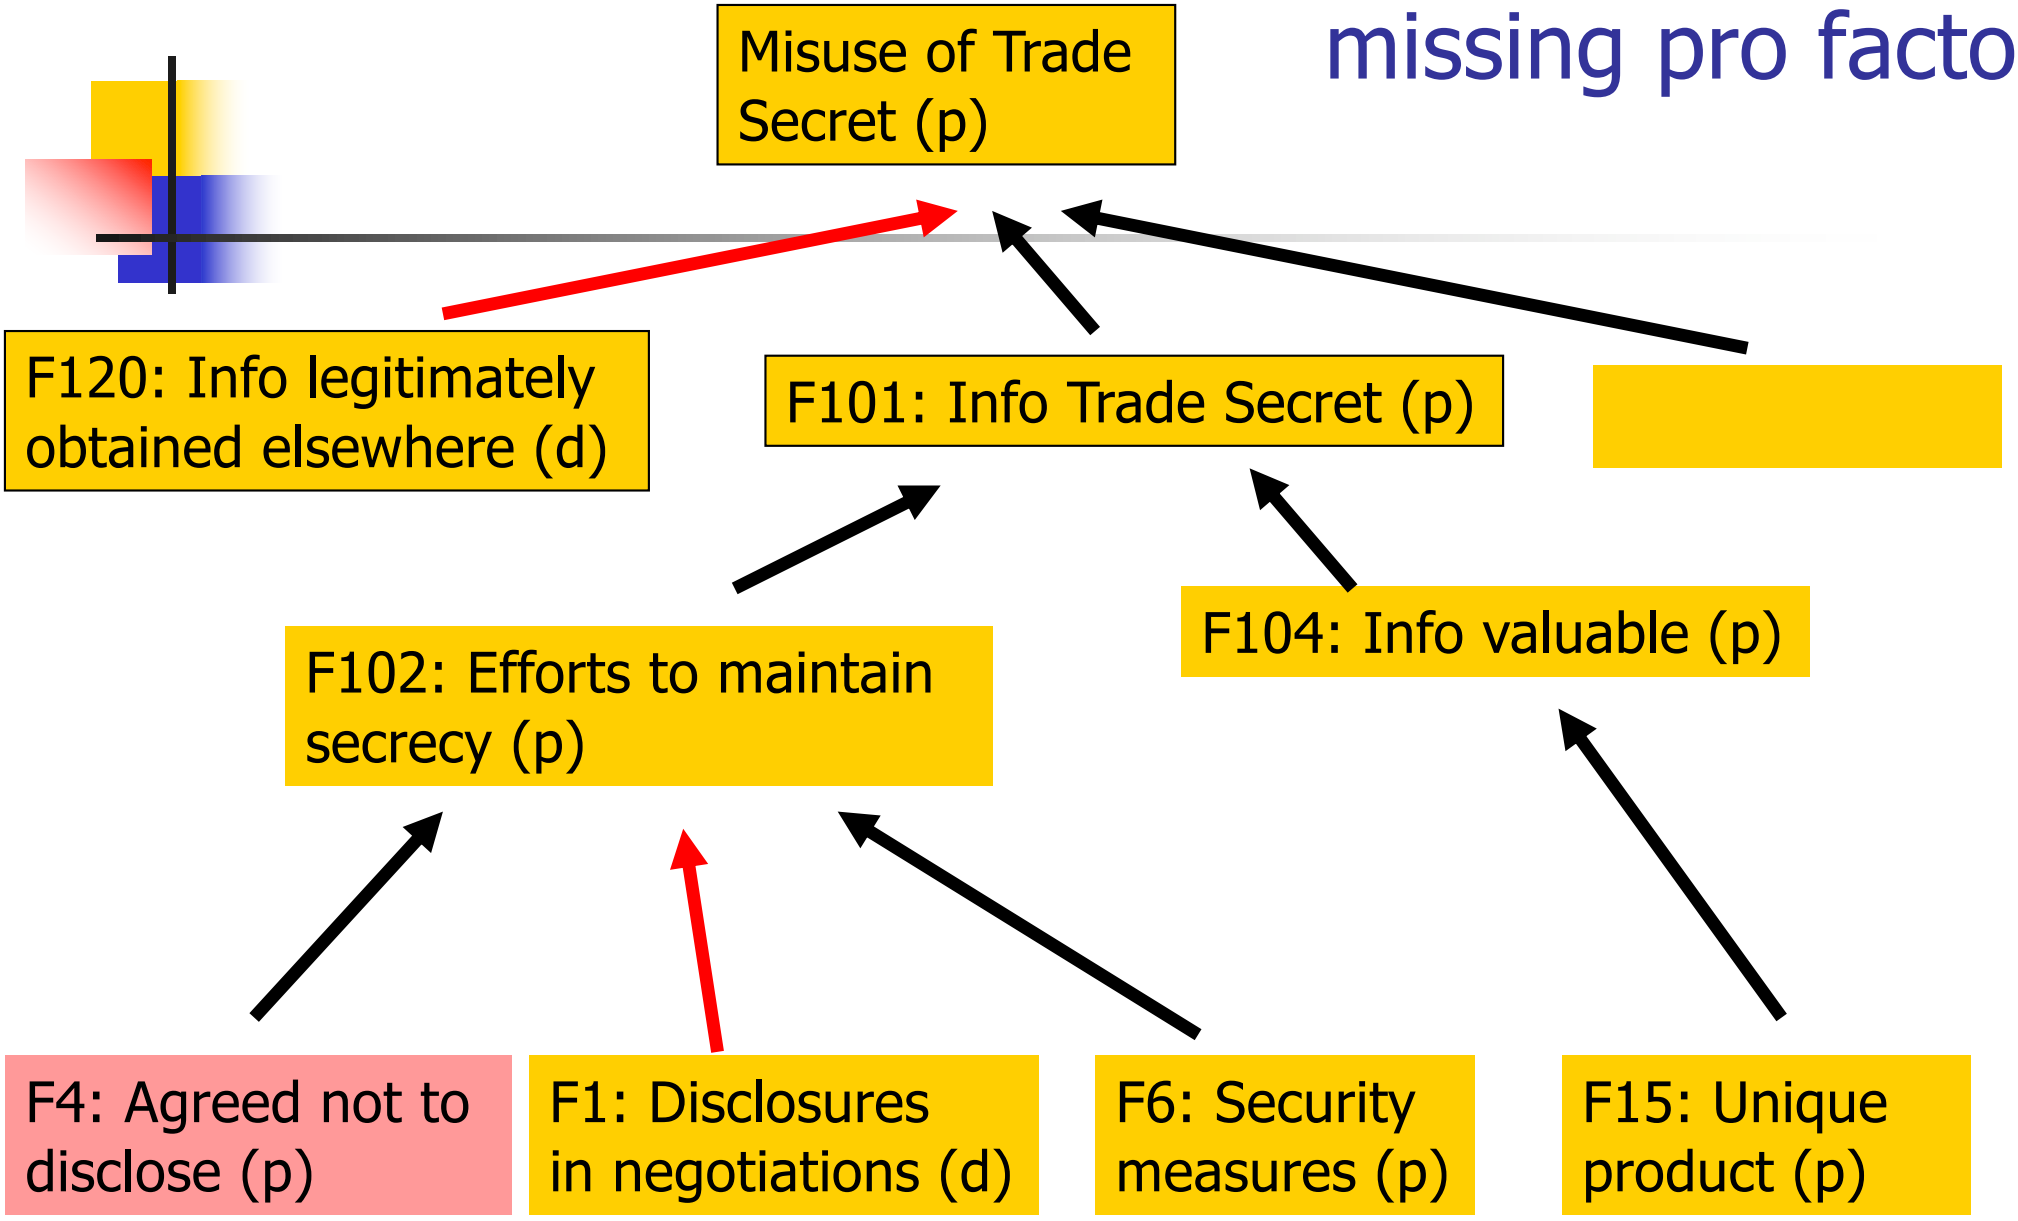

Emphasising distinction

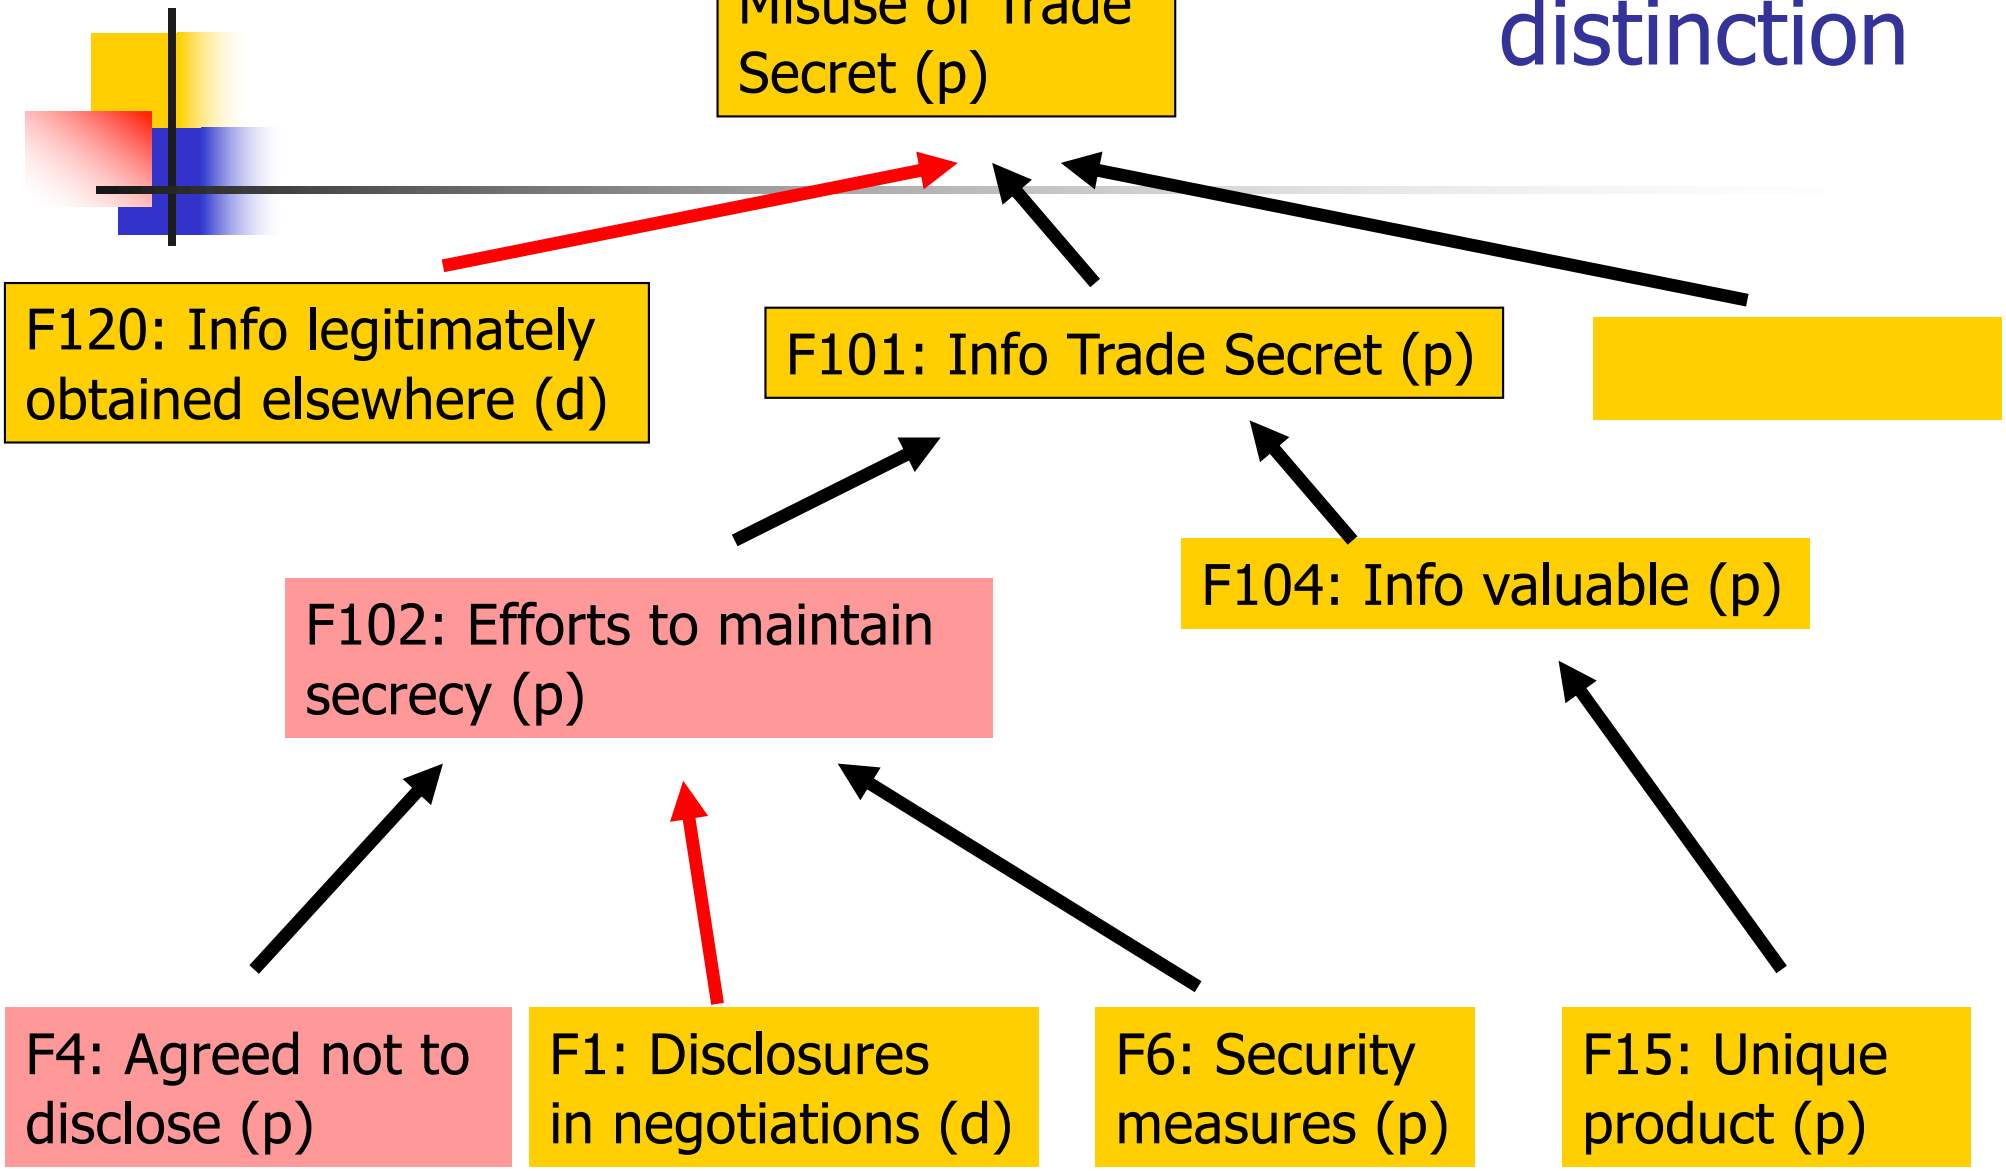

Downplaying by substitution

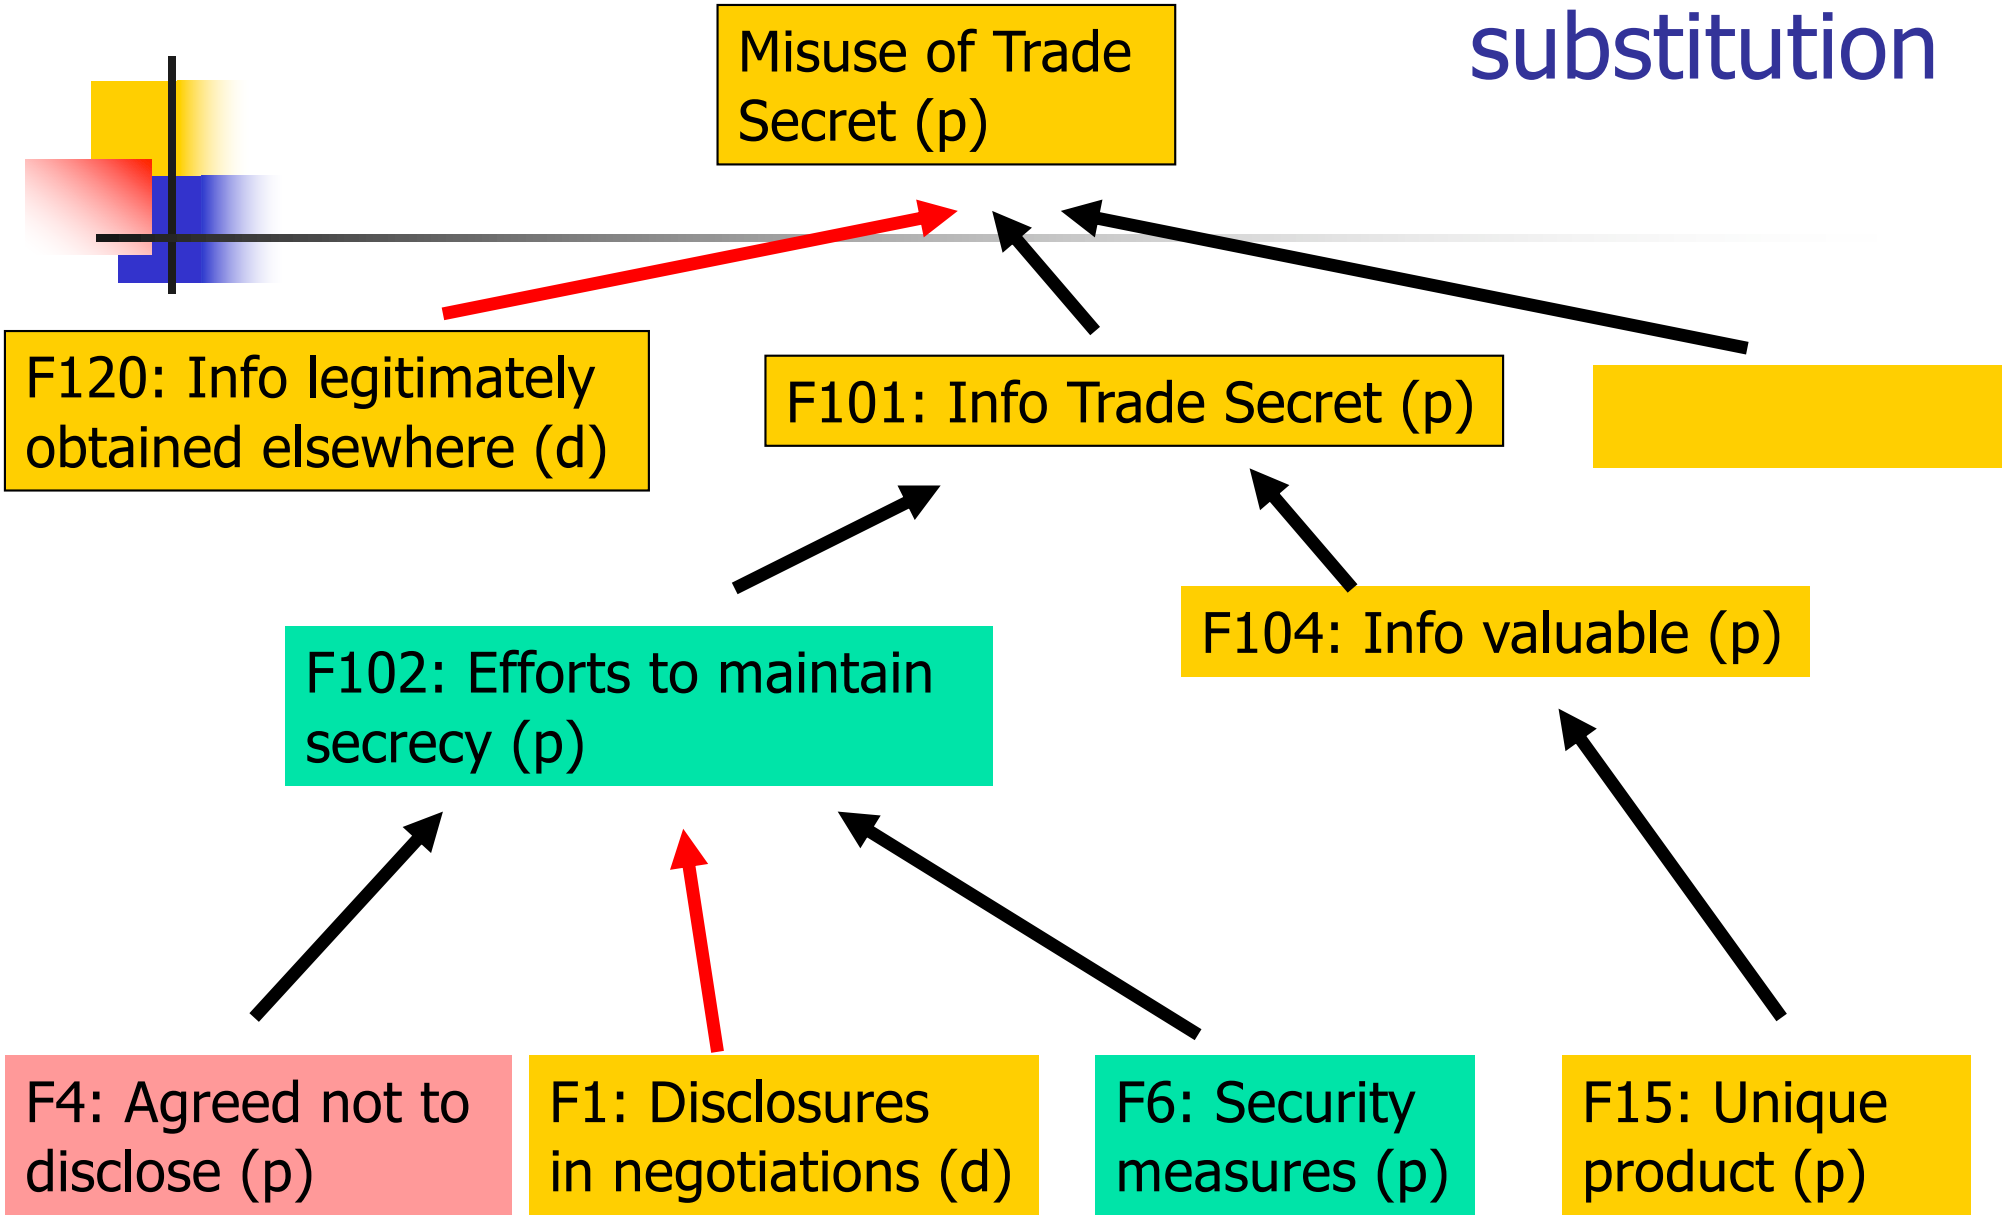

Distinguishing on new con factor

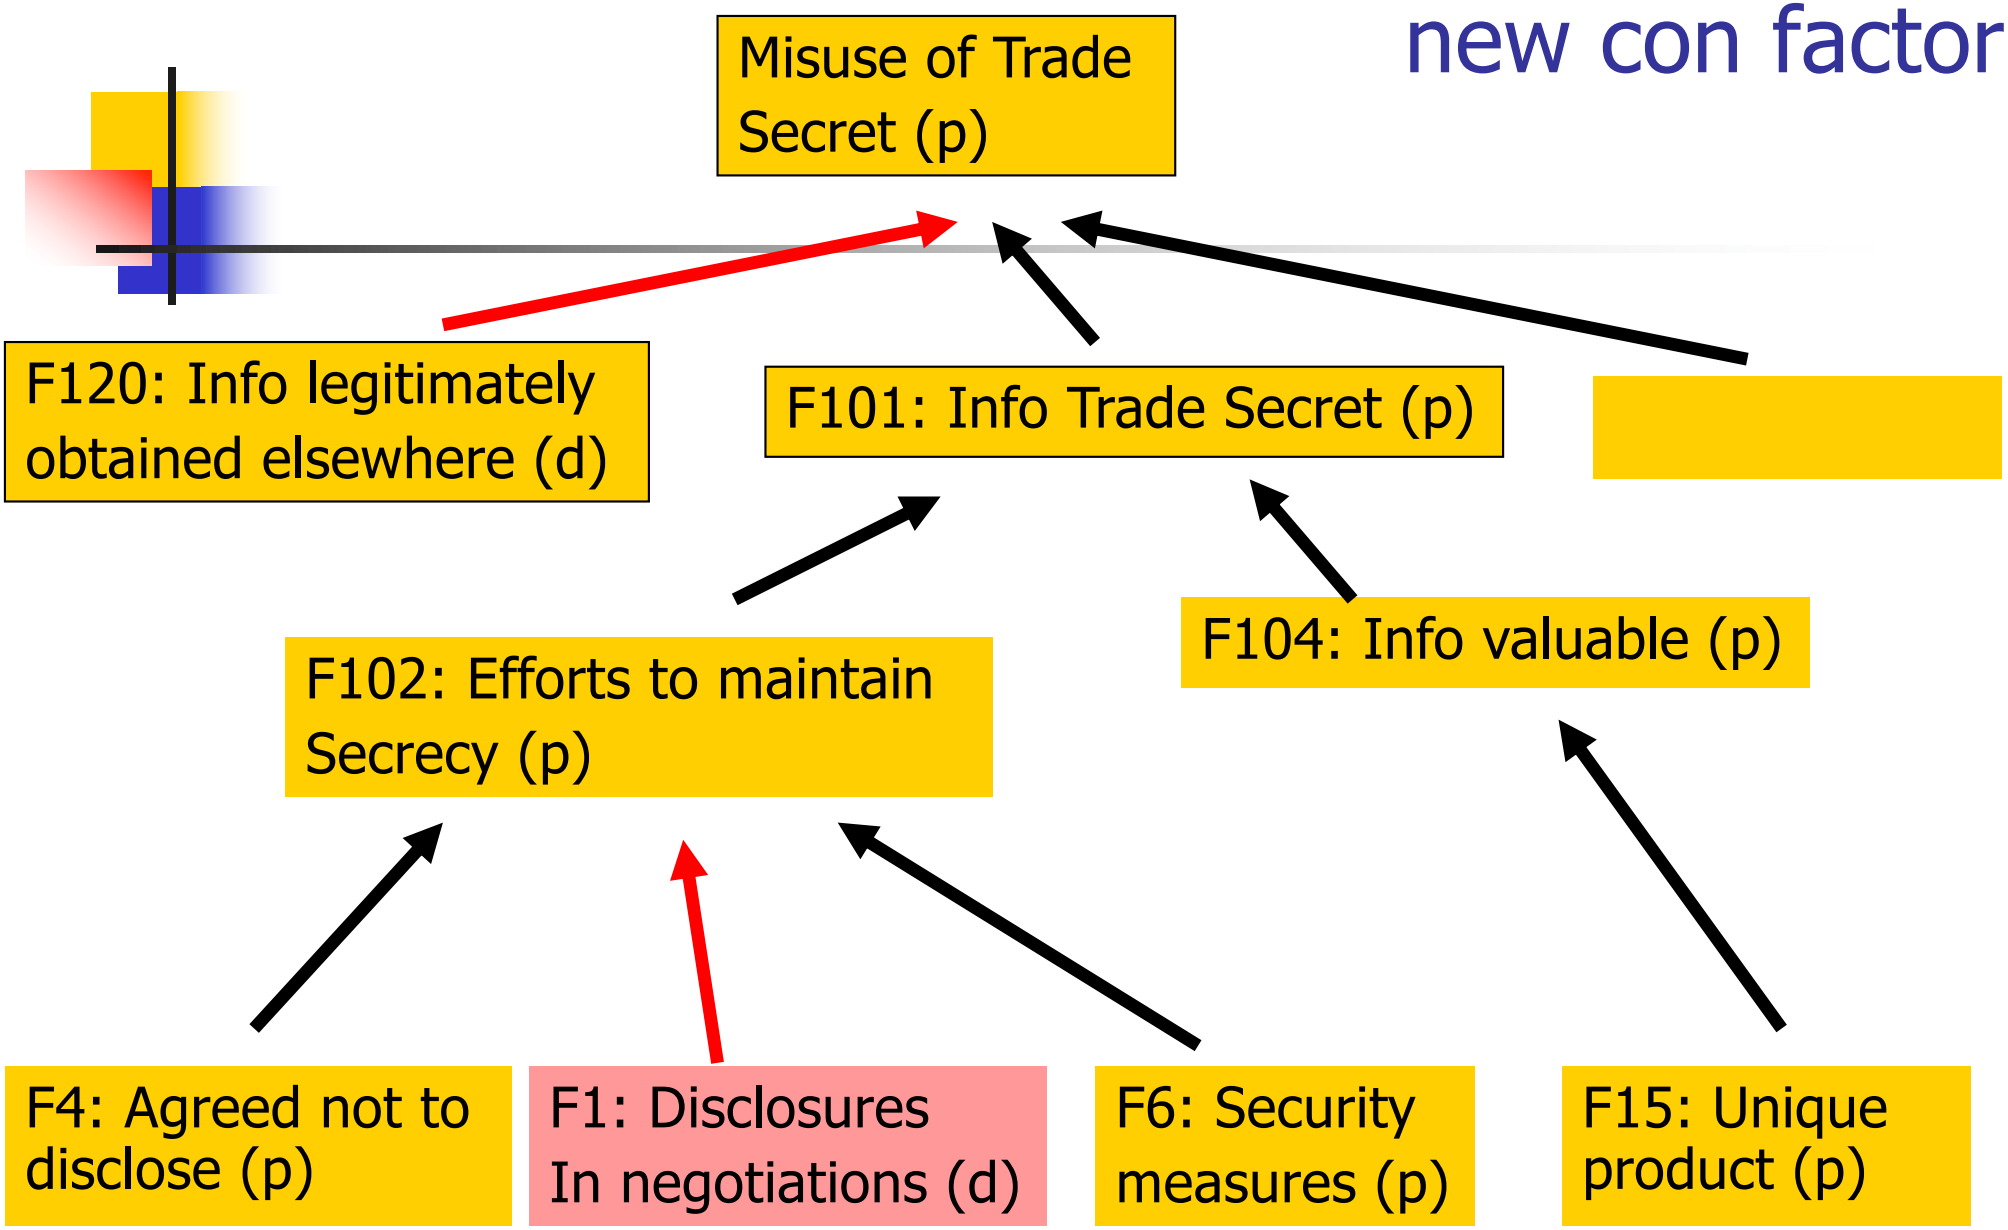

Emphasising distinctions

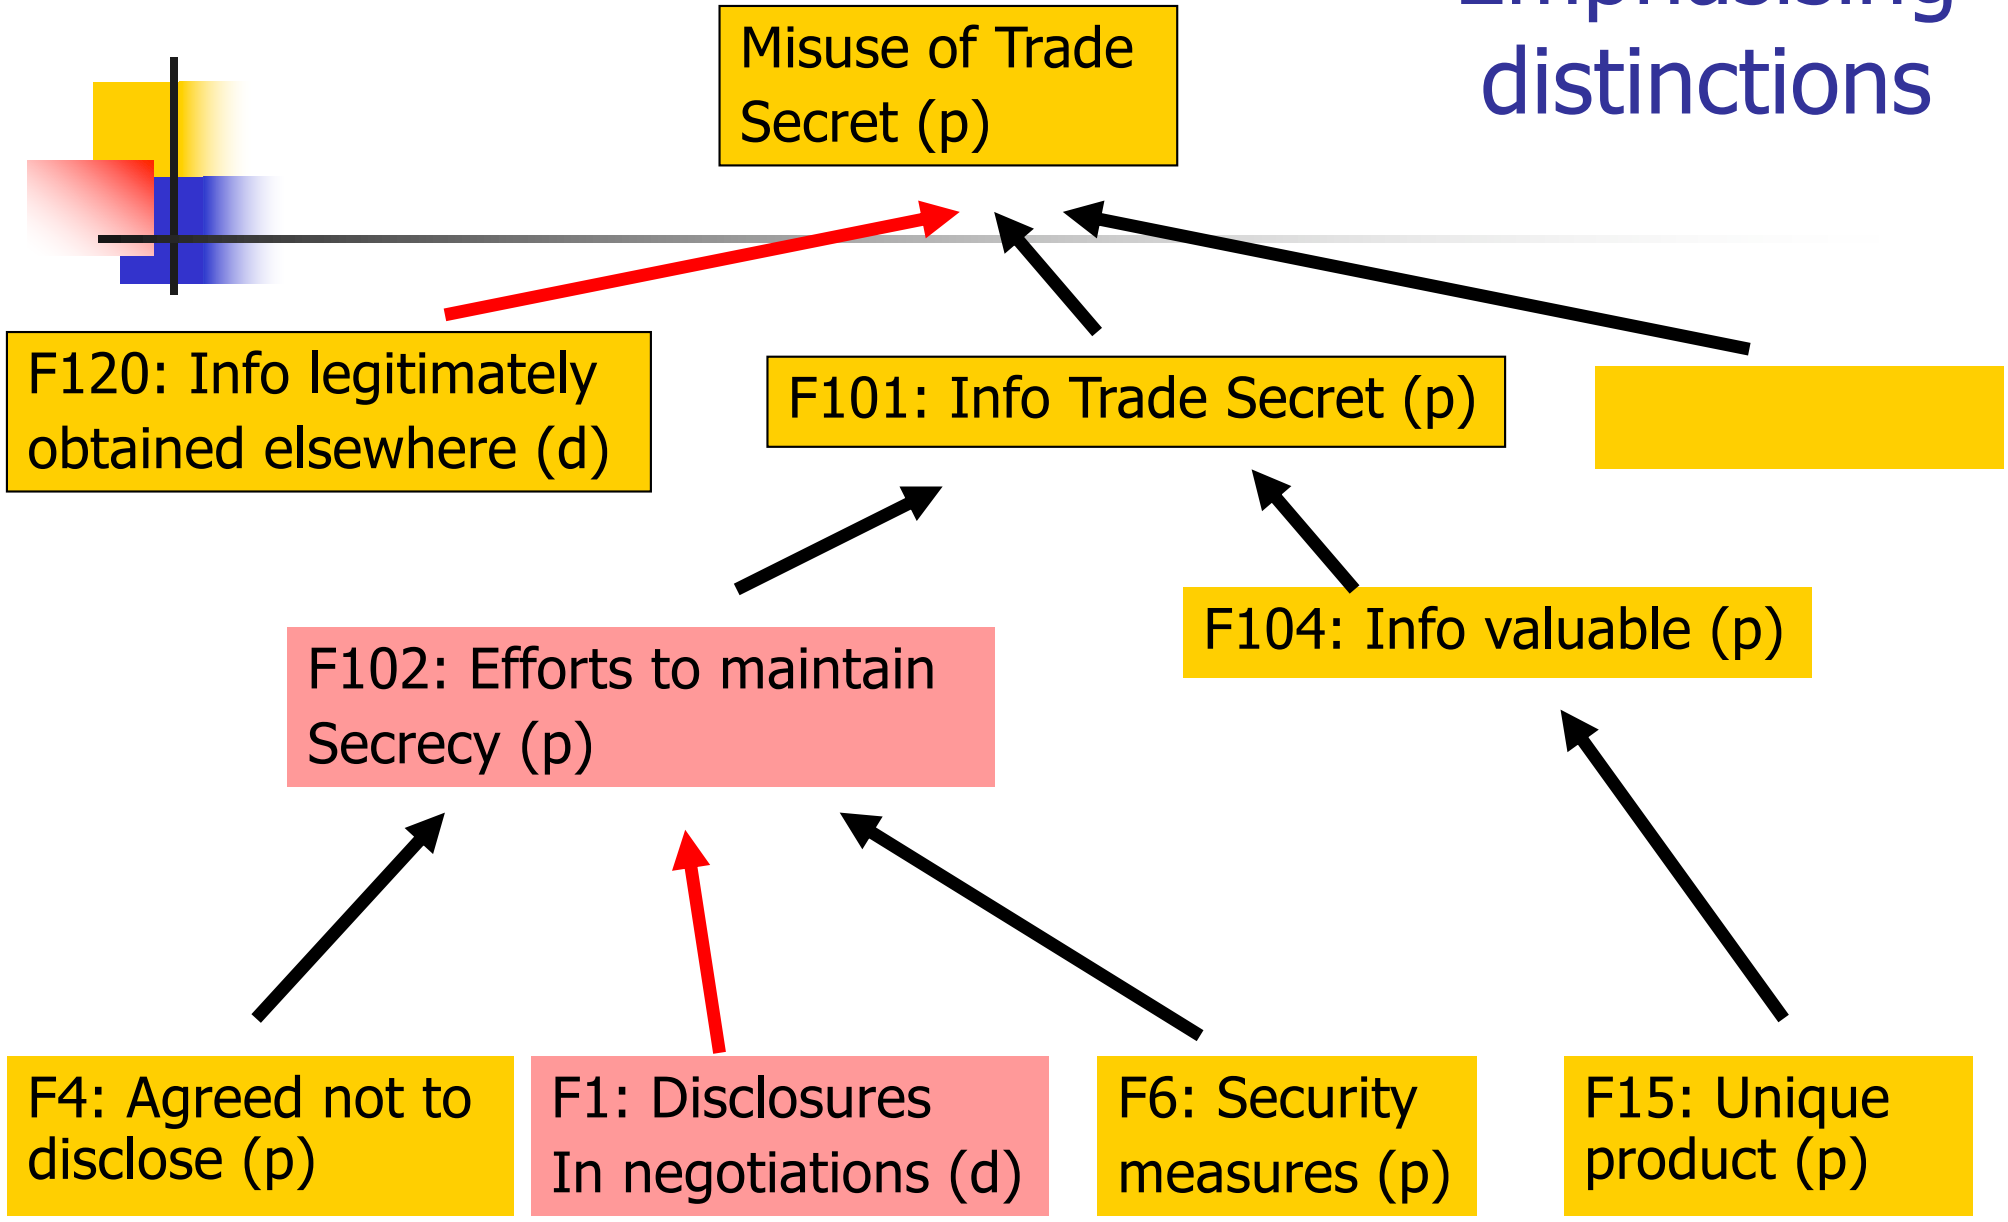

Downplaying by cancellation

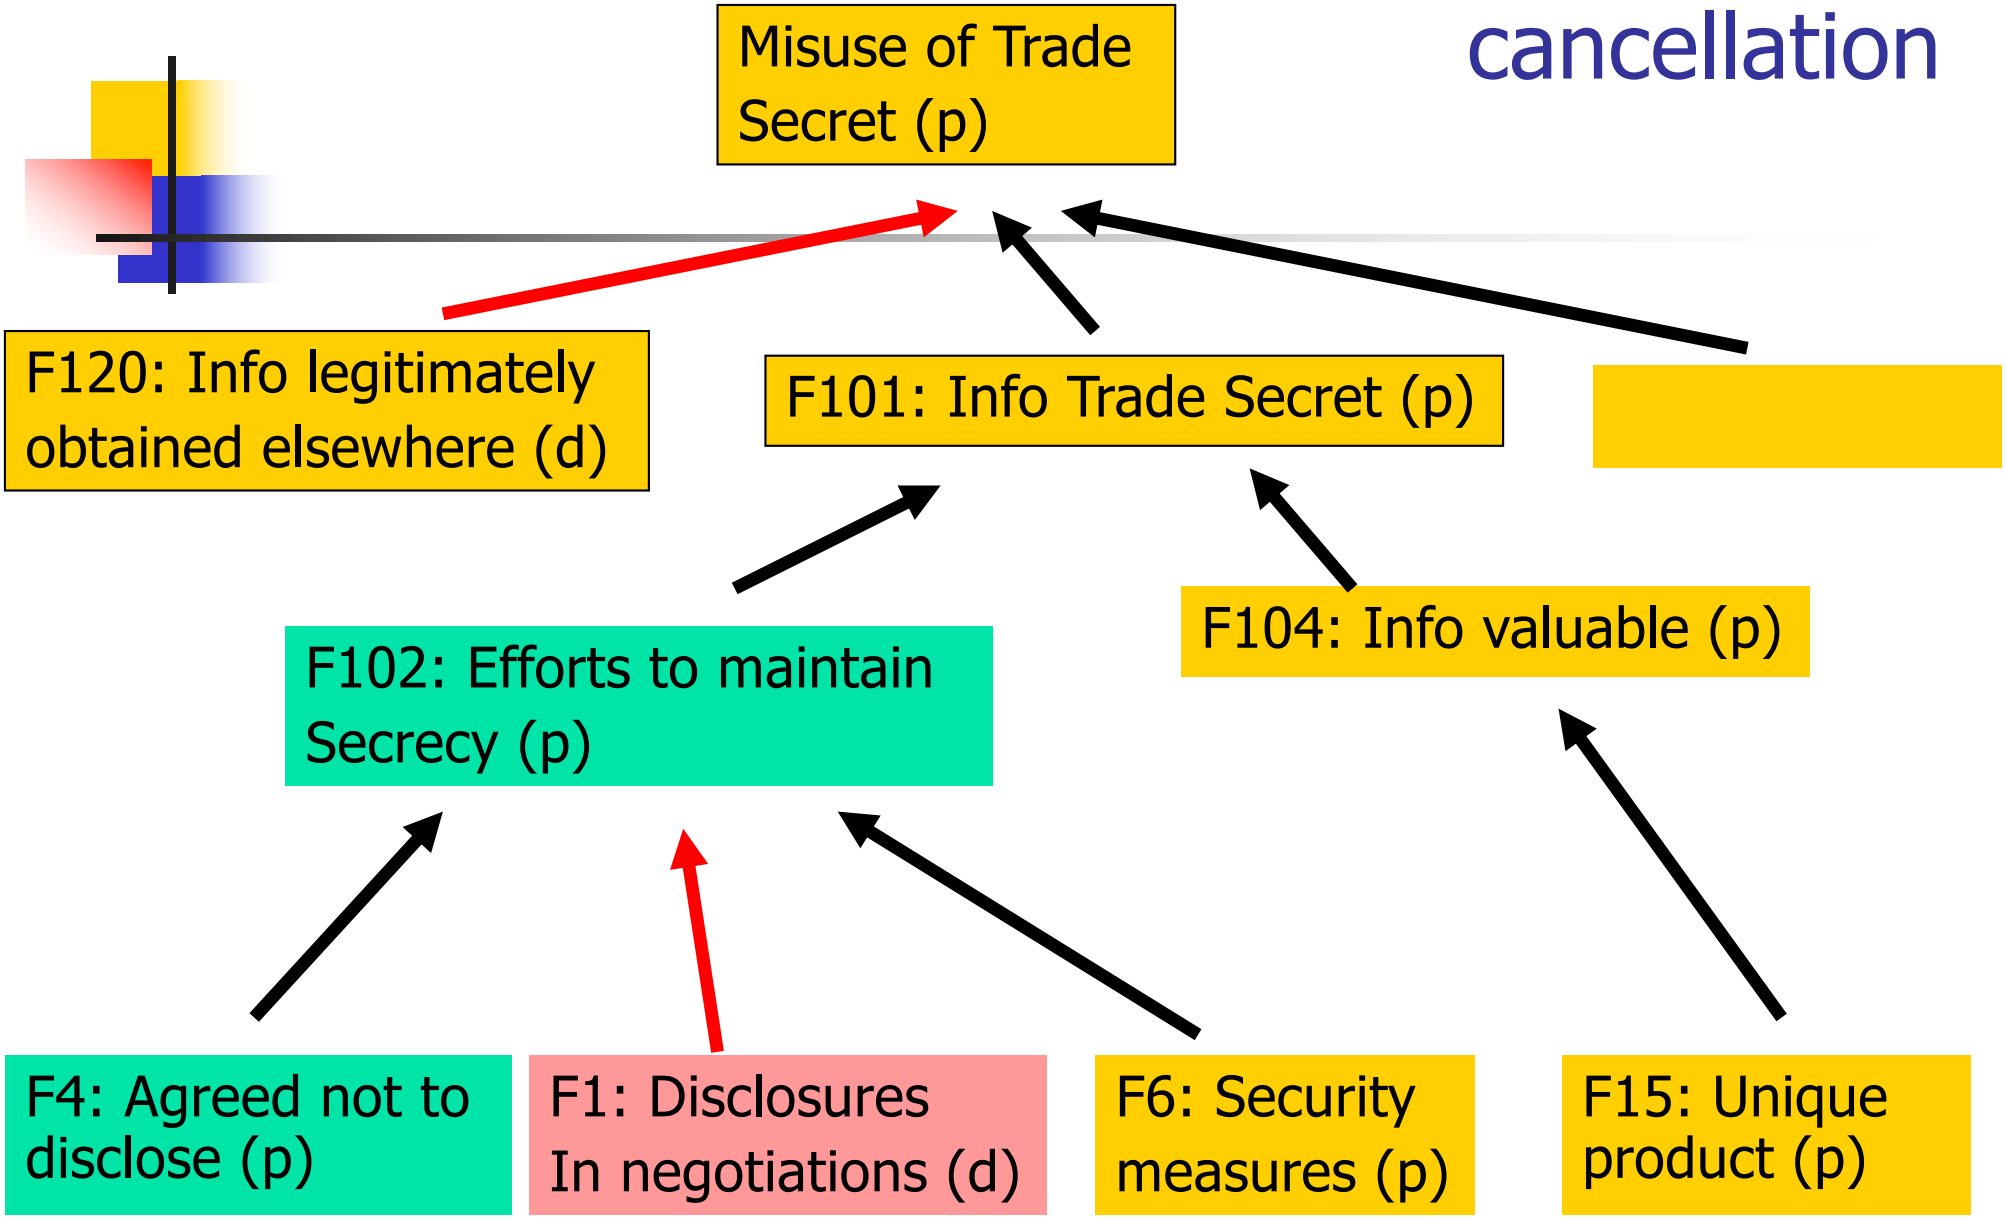

Later developments

- Combining rules and cases
  - Cabaret (Rissland & Skalak 1989-1991)
  - IBP (Ashley & Brüninghaus 2003-2009)
- Underlying values
  - many since Berman & Hafner 1993, Bench-Capon 2000
- Theory construction
  - Bench-Capon & Sartor 2001, Bench-Capon & Chorley 2003-2006
- Precedential constraint
  - Horty 2011-, …

Trevor Bench-Capon

John Horty

Precedential constraint

- When is a decision in a new case 'forced' by a case base?
  - If the case base contains a precedent for that decision that cannot be distinguished:
    - All differences make the new case even stronger for the new decision

J. Horty, Rules and reasons in the theory of precedent. *Legal Theory* 17 (2011): 1-33.

…

H. Prakken, A formal analysis of some factor- and dimension-based accounts of precedential constraint. *Artificial Intelligence and Law* 29 (2021): 559-585.

…

the factors in G favouring s (the pro s factors)

Result model with factors

- For two fact situations F and G and sides s,-s:
- $G \leq_s F$ iff $G^s \subseteq F^s$ and $F^{-s} \subseteq G^{-s}$

  - F is at least as good for s as G iff:
    - F has at least all pro s factors that G has, and
    - F has at most all con s factors that G has
      - Con s = pro -s

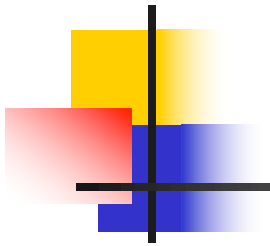

Result model with factors

- Deciding fact situation F for s is forced iff there exists a precedent (G,s) such that $G \leq_s F$
  - Such that F is at least as good for s as G

- Deciding F for s is allowed iff deciding F for -s is not forced.

$X \leq_s Y$ iff $X^s \subseteq Y^s$ and $Y^{-s} \subseteq X^{-s}$ Deciding F for s is forced iff there exists a precedent (X,Y,s) such that X U Y $\leq_s$ F

- *Case 1* – won by plaintiff — $\{\pi 1, \pi 3, \delta 1, \delta 3\} \leq_\pi$ FS1?
  - Deceive-plaintiff ($\pi 1$)
  - Security measures ($\pi 3$)
  - No-unique-product ($\delta 1$)
  - Secrets-disclosed-outsiders ($\delta 3$)
- *Case 2* – won by defendant — $\{\pi 2, \delta 1, \delta 3\} \leq_\delta$ FS1?
  - Bribe-Employee ($\pi 2$)
  - No-unique-product ($\delta 1$)
  - Secrets-disclosed-outsiders ($\delta 3$)
- *FS 1* – $\{\pi 2, \pi 3, \delta 2, \delta 3\}$
  - Bribe-Employee ($\pi 2$)
  - Security measures ($\pi 3$)
  - Info-Reverse-Engineerable ($\delta 2$)
  - Secrets-disclosed-outsiders ($\delta 3$)

$X \leq_s Y$ iff $X^s \subseteq Y^s$ and $Y^{-s} \subseteq X^{-s}$ Deciding F for s is forced iff there exists a precedent (X,Y,s) such that X U Y $\leq_s$ F

- *Case 1* – won by plaintiff — $\{\pi1, \pi3, \delta1, \delta3\} \leq_\pi$ FS1? No
  - Deceive-plaintiff ($\pi1$)
  - Security measures ($\pi3$)
  - No-unique-product ($\delta1$)
  - Secrets-disclosed-outsiders ($\delta3$)
- *Case 2* – won by defendant — $\{\pi2, \delta1, \delta3\} \leq_\delta$ FS1?
  - Bribe-Employee ($\pi2$)
  - No-unique-product ($\delta1$)
  - Secrets-disclosed-outsiders ($\delta3$)
- *FS 1* – — $\{\pi2, \pi3, \delta2, \delta3\}$
  - Bribe-Employee ($\pi2$)
  - Security measures ($\pi3$)
  - Info-Reverse-Engineerable ($\delta2$)
  - Secrets-disclosed-outsiders ($\delta3$)

$X \leq_s Y$ iff $X^s \subseteq Y^s$ and $Y^{-s} \subseteq X^{-s}$ Deciding F for s is forced iff there exists a precedent (X,Y,s) such that X U Y $\leq_s$ F

- *Case 1* – won by plaintiff — $\{\pi 1, \pi 3, \delta 1, \delta 3\} \leq_\pi$ FS1?
  - Deceive-plaintiff ($\pi 1$)
  - Security measures ($\pi 3$)
  - No-unique-product ($\delta 1$)
  - Secrets-disclosed-outsiders ($\delta 3$)
- *Case 2* – won by defendant — $\{\pi 2, \delta 1, \delta 3\} \leq_\delta$ FS1?
  - Bribe-Employee ($\pi 2$)
  - No-unique-product ($\delta 1$)
  - Secrets-disclosed-outsiders ($\delta 3$)
- *FS 1* – $\{\pi 2, \pi 3, \delta 2, \delta 3\}$
  - Bribe-Employee ($\pi 2$)
  - Security measures ($\pi 3$)
  - Info-Reverse-Engineerable ($\delta 2$)
  - Secrets-disclosed-outsiders ($\delta 3$)

$X \leq_s Y$ iff $X^s \subseteq Y^s$ and $Y^{-s} \subseteq X^{-s}$
Deciding F for s is forced iff there exists a precedent (X,Y,s) such that X U Y $\leq_s$ F

- *Case 1* – won by plaintiff — $\{\pi 1, \pi 3, \delta 1, \delta 3\} \leq_\pi$ FS1?
  - Deceive-plaintiff ($\pi 1$)
  - Security measures ($\pi 3$)
  - No-unique-product ($\delta 1$)
  - Secrets-disclosed-outsiders ($\delta 3$)
- *Case 2* – won by defendant — $\{\pi 2, \delta 1, \delta 3\} \leq_\delta$ FS1? No
  - Bribe-Employee ($\pi 2$)
  - No-unique-product ($\delta 1$)
  - Secrets-disclosed-outsiders ($\delta 3$)
- *FS 1* – — $\{\pi 2, \pi 3, \delta 2, \delta 3\}$
  - Bribe-Employee ($\pi 2$)
  - Security measures ($\pi 3$)
  - Info-Reverse-Engineerable ($\delta 2$)
  - Secrets-disclosed-outsiders ($\delta 3$)

Example dimensions in HYPO

- Number of outsider disclosees (0,1,….)

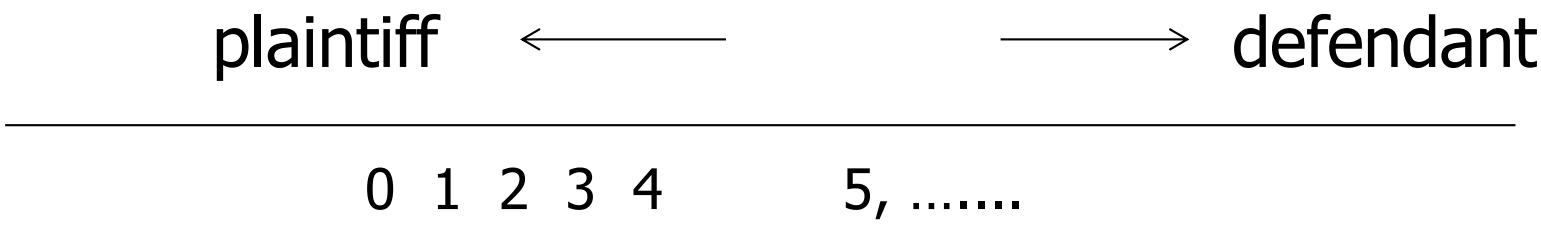

Example dimensions in HYPO

- Security measures (minimal measures, acces to premises controlled, entry by visitors restricted, entry by employees restricted)

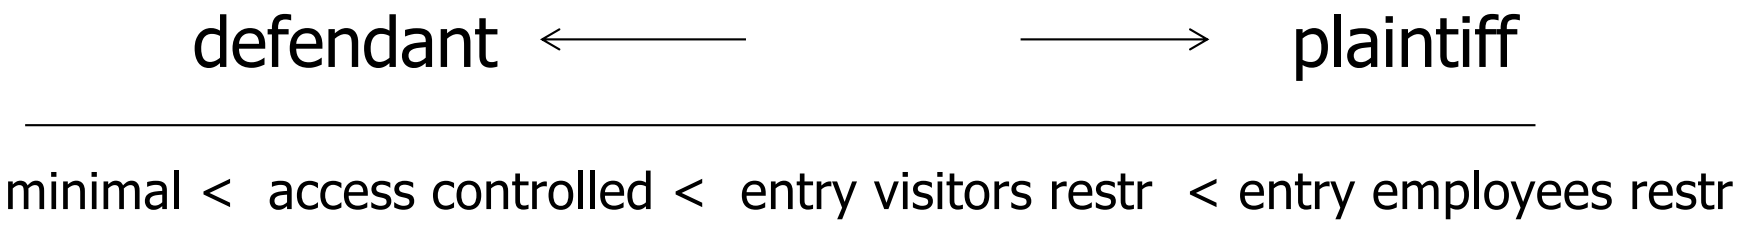

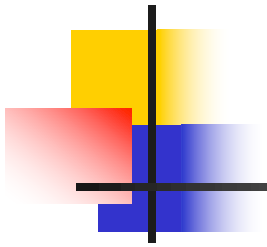

Result model with dimensions

- For any two fact situations F and G given a set of dimensions:
  - $G \leq_s F$ iff F is for every dimension at least good for s as G.
- Deciding fact situation F for s is forced iff there exists a precedent with fact situation G and decided for s such that $G \leq_s F$

Example (result model for dimensions)

- *Case 1* – won by plaintiff
  - d1: Deceive-plaintiff = yes
  - d2: Bribe-Employee = no
  - d3: Security measures = Entry-By-Visitors-Restricted
  - d4: Unique-product = no
  - d5: Info-Reverse-Engineerable = no
  - d6: Secrets-disclosed-outsiders = 20

- *Case 2* –
  - d1: Deceive-plaintiff = no
  - d2: Bribe-Employee = yes
  - d3: Security measures = Minimal
  - d4: Unique-product = no
  - d5: Info-Reverse-Engineerable = yes
  - d6: Secrets-disclosed-outsiders = 0

d1: Deceive-plaintiff = {no $<_\pi$ yes}
d2: Bribe-Employee = {no $<_\pi$ yes}
d3: Security-Measures = {Minimal-Measures $<_\pi$
Access-To-Premises-Controlled $<_\pi$
Entry-By-Visitors-Restricted $<_\pi$
Restrictions-On-Entry-By-Employees}
d4: Unique-product = {yes $<_\delta$ no}
d5: Info-Reverse-Engineerable = {no $<_\delta$ yes}
d6: Secret-disclosed-outsiders = {0 $<_\delta$ 1 $<_\delta$ 2 $<_\delta$ …}

Is Case 2 forced for plaintiff?

Example (result model for dimensions)

- *Case 1* – won by plaintiff
  - d1: Deceive-plaintiff = yes
  - d2: Bribe-Employee = no
  - d3: Security measures = Entry-By-Visitors-Restricted
  - d4: Unique-product = no
  - d5: Info-Reverse-Engineerable = no
  - d6: Secrets-disclosed-outsiders = 20
- *Case 2* –
  - d1: Deceive-plaintiff = no
  - d2: Bribe-Employee = yes
  - d3: Security measures = Minimal
  - d4: Unique-product = no
  - d5: Info-Reverse-Engineerable = yes
  - d6: Secrets-disclosed-outsiders = 0

d1: Deceive-plaintiff = {no $<_\pi$ yes}
d2: Bribe-Employee = {no $<_\pi$ yes}
d3: Security-Measures = {Minimal-Measures $<_\pi$
Access-To-Premises-Controlled $<_\pi$
Entry-By-Visitors-Restricted $<_\pi$
Restrictions-On-Entry-By-Employees}
d4: Unique-product = {yes $<_\delta$ no}
d5: Info-Reverse-Engineerable = {no $<_\delta$ yes}
d6: Secret-disclosed-outsiders = {0 $<_\delta$ 1 $<_\delta$ 2 $<_\delta$ ...}

Is Case 2 forced for plaintiff? No.

Rosa Ratsma

‘Explaining’ ML predictions with CBR

- Regard training data as cases/precedents
  - Direction of factors/dimensions can be learned from the BB
- Use AI & law model of CBR for explaining a prediction of the BB model
  - If a precedent forces the prediction, then show it
  - Otherwise explain away the relevant differences

- H. Prakken & R. Ratsma, A top-level model of case-based argumentation for explanation: formalisation and experiments. *Argument and Computation* 13 (2022): 159-194
- W. van Woerkom et al. A Fortiori case-based reasoning: from theory to data. *JAIR* 81 (2024): 401-441.
- J. Peters et al., Model- and data-agnostic justifications with a fortiori case-based argumentation. *Proceedings of ICAIL 2025, pp. 102-111.*

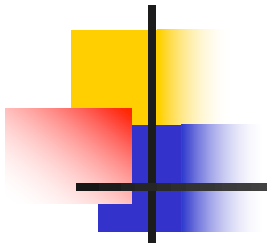

Example (churning)

- *John will stay with us, since like Maria he recently visited our website. (citation)*
- *But Maria also downloaded software from our site while John didn't (missing pro). Moreover, John did not reply to our last mailing (new con)*
- *But John downloaded a brochure, so like Maria he showed an interest in our products (substitution). Moreover, John contacted our help desk, so like Maria he remained in contact with us (cancellation)*

Wijnand van Woerkom

Follow-up work on precedential constraint

- Discovering dimension orderings in datasets
- Intermediate factors in result model of precedential constraint
  - Assuming a CATO-style factor hierarchy
  - Generalising it to dimensions
- Implementations with SAT solver

- W. van Woerkom et al. A Fortiori Case-Based Reasoning: From Theory to Data. *JAIR* 81 (2024): 401-441.
- W. van Woerkom et al., Hierarchical models of precedential constraint. *Artificial Intelligence and Law*, in press.

Hierarchical precedential constraint

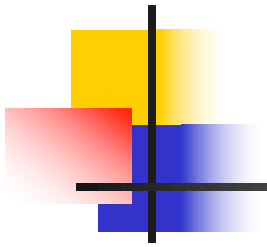

- Apply precedential constraint to
  - decide on intermediate factors
  - decide on the final outcome given intermediate factors
    - With dimensions: allowed or forced ranges of values

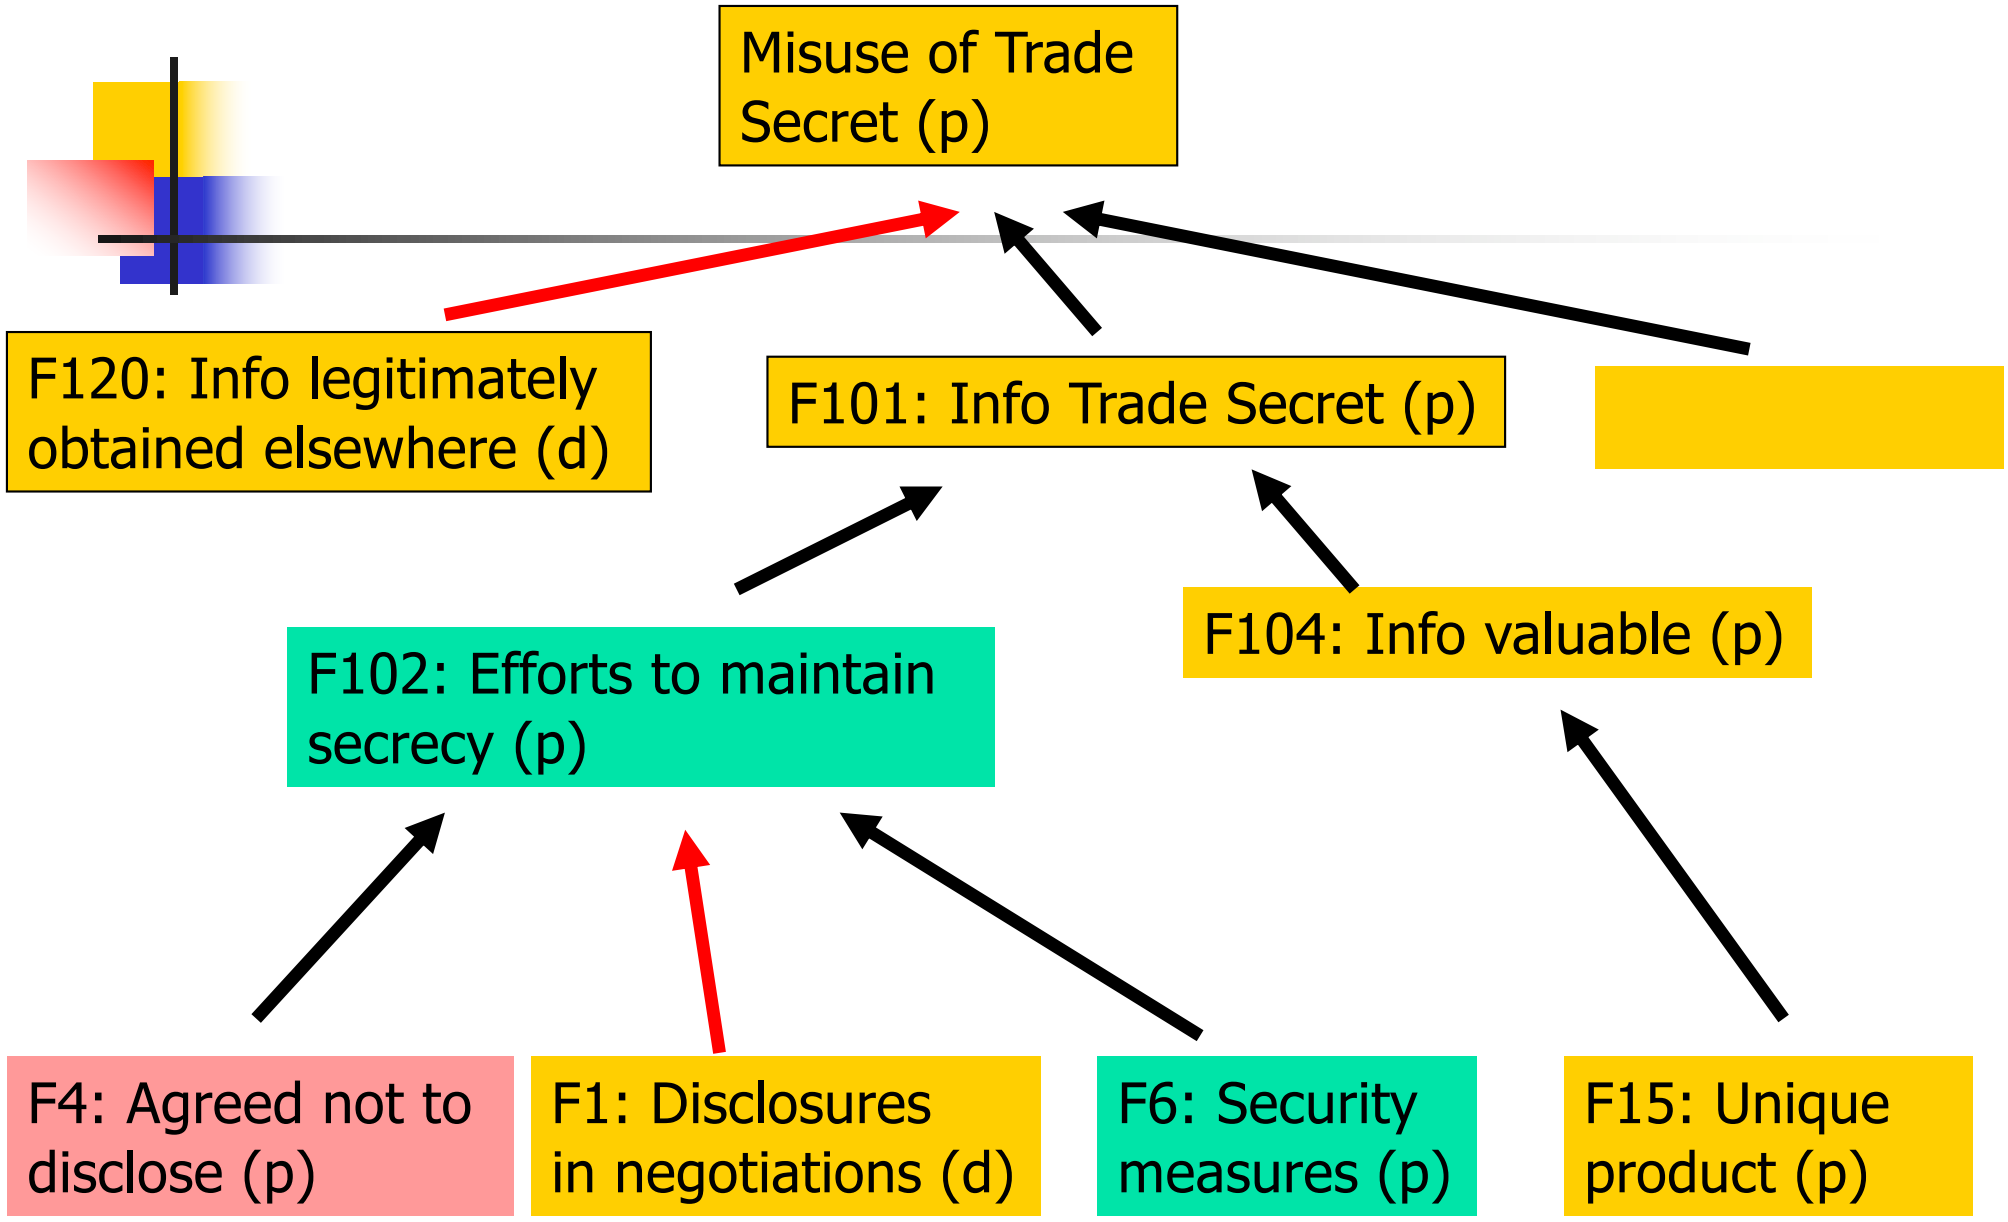
Misuse of Trade Secret (p)
F120: Info legitimately obtained elsewhere (d)
F101: Info Trade Secret (p)
F102: Efforts to maintain secrecy (p)
F104: Info valuable (p)
F4: Agreed not to disclose (p)
F1: Disclosures in negotiations (d)
F6: Security measures (p)
F15: Unique product (p)

Joeri Peters

Consistency of datasets

- A case base is inconsistent iff it forces opposite outcomes for the same fact situation
- Degree of consistency of a CB:
  - The rate of pairs (F,s), (G,-s) for which $F \leq_s G$
- Authoritativeness of a case (F,s): the proportion of cases (G,s) and (G',-s) such that $F \leq_s G$ and $F \leq_s G'$

- R. Ratsma (2020), *Unboxing the Black Box using Case-based Argumentation*. AI Msc thesis, Utrecht University
- J. Peters et al., Justifications derived from inconsistent case bases using authoritativeness. *Proceedings of ArgXAI 2022.*
- *W. van Woerkom,* Formal results on case-base consistency: a COMPAS case study. *Proceedings of ICAIL*

Decision support for deciding on fitness to drive

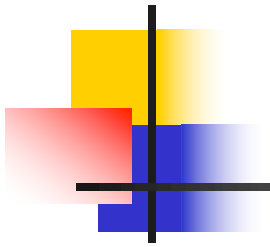

- At Dutch Central Driving License Office
  - "No machine learning"
    - Not transparent or explainable
  - "Is standard CBR useful?"
- Standard CBR:
  - assumes 'features' without preferences for decisions
  - applies numerical similarity measures
  - suggests decision of precedent(s) with highest similarity to current case

Joep Nouwens: Msc project AI-UU

- Combine and compare with AI & law style CBR:
  - Features with preferences for decisions
    - Not all differences are relevant!
  - Decision-making by precedential constraint
- Example dimensions
  - Heart disease?
  - Bipolar disorder?
  - Eye sight
  - Epiliptic attacks
  - …

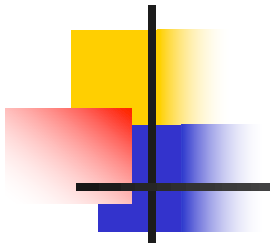

Experiment

- Case base: 15.843 cases, 123 dimensions
  - 80% used as precedent
  - 20% used as test case
- Four decision rules:
  - Standard CBR: predict decision with highest similarity
  - Precedential constraint with if both allowed/forced:
    - predict 'fit'
    - predict 'unfit'
    - predict decision of case with highest similarity according to standard CBR

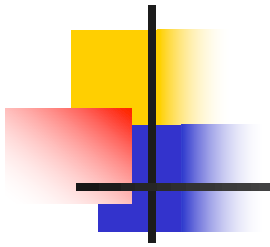

Experiments with accuracies

- Case base: 15.843 cases, 123 dimensions
  - 80% used as precedent
  - 20% used as test case
- Four decision rules:
  - Standard CBR: predict decision with highest similarity (92%)
  - Precedential constraint with if both allowed/forced:
    - predict 'fit' (70%)
    - predict 'unfit' (64%)
    - predict decision of case with highest similarity according to standard CBR (91%)

The value of predictive experiments

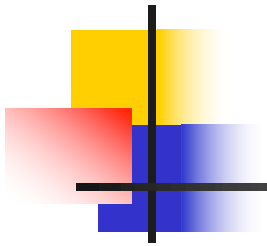

- "High predictive accuracy is evidence of legal correctness of the model"
  - Aleven, Ashley
- HP: only true if system and humans:
  - apply the same knowledge
  - reason with it in the same way

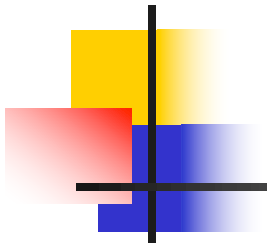

Conclusions

- There is interesting work on symbolic models of legal case-based reasoning
- With recent research applications in
  - Explainable AI
  - Analysing consistency of decision-making

- Initial work on letting LLMs
  - provide the inputs of symbolic models
  - generate case-based arguments
  - (Ashley & students since 2023)