Privacy Paradox: When Does Hiding in Plain Sight Work?

Tatiana Mayskaya[1] Arina Nikandrova[2]

[1]Higher School of Economics

[2]City, University of London

HSE
3 June 2021

Prince Harry and Meghan Markle: In pursuit of privacy

Prince Harry's lawsuit against tabloids could backfire, commentators claim

Duke of Sussex's legal action against Sun and Daily Mirror over alleged phone hacking takes attack on press up a level

The Guardian, 5 October 2019

Prince Harry and Meghan Markle: In pursuit of privacy

Prince Harry's lawsuit against tabloids could backfire, commentators claim

Duke of Sussex's legal action against Sun and Daily Mirror over alleged phone hacking takes attack on press up a level

The Guardian, 5 October 2019

The Sussexes have reportedly collaborated with a book about their split from the royal family. So much for pursuing a new life of privacy

The Guardian, 5 May 2020

Framework outline

- Two players: a hider (she) and a seeker (he)

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories
- Hider's story is protected at some exogenous default level
  - default protection: privacy laws, encryption of data in WhatsApp, US tradition of no trespassing, royal family security protocols

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories
- Hider's non-compromising story is protected at some exogenous default level
  - default protection: privacy laws, encryption of data in WhatsApp, US tradition of no trespassing, royal family security protocols

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories
- Hider's non-compromising story is protected at some exogenous default level
  - default protection: privacy laws, encryption of data in WhatsApp, US tradition of no trespassing, royal family security protocols
- Hider chooses the level of protection for compromising story
  - selective protection: one's habits on social media, company's press releases

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories
- Hider's non-compromising story is protected at some exogenous default level
  - default protection: privacy laws, encryption of data in WhatsApp, US tradition of no trespassing, royal family security protocols
- Hider chooses the level of protection for compromising story
  - selective protection: one's habits on social media, company's press releases
  - selective protection is costless but subject to a constraint

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories
- Hider's non-compromising story is protected at some exogenous default level
  - default protection: privacy laws, encryption of data in WhatsApp, US tradition of no trespassing, royal family security protocols
- Hider chooses the level of protection for compromising story
  - selective protection: one's habits on social media, company's press releases
  - selective protection is costless but subject to a constraint
- Hider publicly commits to the level of selective protection

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories
- Hider's non-compromising story is protected at some exogenous default level
  - default protection: privacy laws, encryption of data in WhatsApp, US tradition of no trespassing, royal family security protocols
- Hider chooses the level of protection for compromising story
  - selective protection: one's habits on social media, company's press releases
  - selective protection is costless but subject to a constraint
- Hider publicly commits to the level of selective protection
- The story type and the realised level of protection are initially private to the hider

Framework outline

- Two players: a hider (she) and a seeker (he)
- Hider's story is either compromising or non-compromising
- Seeker aims to find and expose only compromising stories
- Hider's non-compromising story is protected at some exogenous default level
  - default protection: privacy laws, encryption of data in WhatsApp, US tradition of no trespassing, royal family security protocols
- Hider chooses the level of protection for compromising story
  - selective protection: one's habits on social media, company's press releases
  - selective protection is costless but subject to a constraint
- Hider publicly commits to the level of selective protection
- The story type and the realised level of protection are initially private to the hider
- Seeker can learn the story at a cost

Other examples

Other examples

- company hiding financial problems

Other examples

- company hiding financial problems
- bank hiding the depletion of cash reserves

Other examples

- company hiding financial problems
- bank hiding the depletion of cash reserves
- politician hiding her misdeeds

Results preview

- What is the optimal level of selective protection?

Results preview

- What is the optimal level of selective protection?
- Trade-off: stronger protection $\downarrow$ effectiveness and $\uparrow$ duration of learning

Results preview

- What is the optimal level of selective protection?
- Trade-off: stronger protection $\downarrow$ effectiveness and $\uparrow$ duration of learning

Result 1

If the default protection is weak, the optimal strength of selective protection is low

Results preview

- What is the optimal level of selective protection?
- Trade-off: stronger protection $\downarrow$ effectiveness and $\uparrow$ duration of learning

Result 1

If the default protection is weak, the optimal strength of selective protection is low

Result 2

Hider benefits from stronger default protection

Results preview

- What is the optimal level of selective protection?
- Trade-off: stronger protection $\downarrow$ effectiveness and $\uparrow$ duration of learning

Result 1

If the default protection is weak, the optimal strength of selective protection is low

Result 2

Hider benefits from stronger default protection

Privacy Paradox *(Norberg et al (2007))*

People say they value privacy highly but behave as if they value it very little.

Results preview

- What is the optimal level of selective protection?
- Trade-off: stronger protection $\downarrow$ effectiveness and $\uparrow$ duration of learning

Result 1

If the default protection is weak, the optimal strength of selective protection is low

Result 2

Hider benefits from stronger default protection

Privacy Paradox *(Norberg et al (2007))*

People say they value privacy highly but behave as if they value it very little.

Results preview

- What is the optimal level of selective protection?
- Trade-off: stronger protection $\downarrow$ effectiveness and $\uparrow$ duration of learning

Result 1

If the default protection is weak, the optimal strength of selective protection is low

Result 2

Hider benefits from stronger default protection

$\Rightarrow$ strict privacy laws

Privacy Paradox *(Norberg et al (2007))*

People say they value privacy highly but behave as if they value it very little.

Results preview

- What is the optimal level of selective protection?
- Trade-off: stronger protection $\downarrow$ effectiveness and $\uparrow$ duration of learning

Result 1

If the default protection is weak, the optimal strength of selective protection is low

Result 2

Hider benefits from stronger default protection
$\Rightarrow$ strict privacy laws

Privacy Paradox *(Norberg et al (2007))*

People say they value privacy highly but behave as if they value it very little.

Result 3

When the hider controls the number of seekers, the open access policy with infinite number of seekers is optimal if the discount rate is sufficiently low.

Outline

Model

- Hider commits to the level of selective protection, parametrized by $\mu_1 > 0$. The level of default protection $\mu_0 > 0$ is exogenous.

Model

hider chooses
selective
protection

Model

- Hider commits to the level of selective protection, parametrized by $\mu_1 > 0$. The level of default protection $\mu_0 > 0$ is exogenous.
- Hider gets involved in a story of type $\theta \in \{0, 1\}$
  - $\theta = 1$: compromising story
  - $\theta = 0$: non-compromising story

  Seeker knows μ_0 and μ_1, and believes $\Pr(\theta = 1) = p$

Model

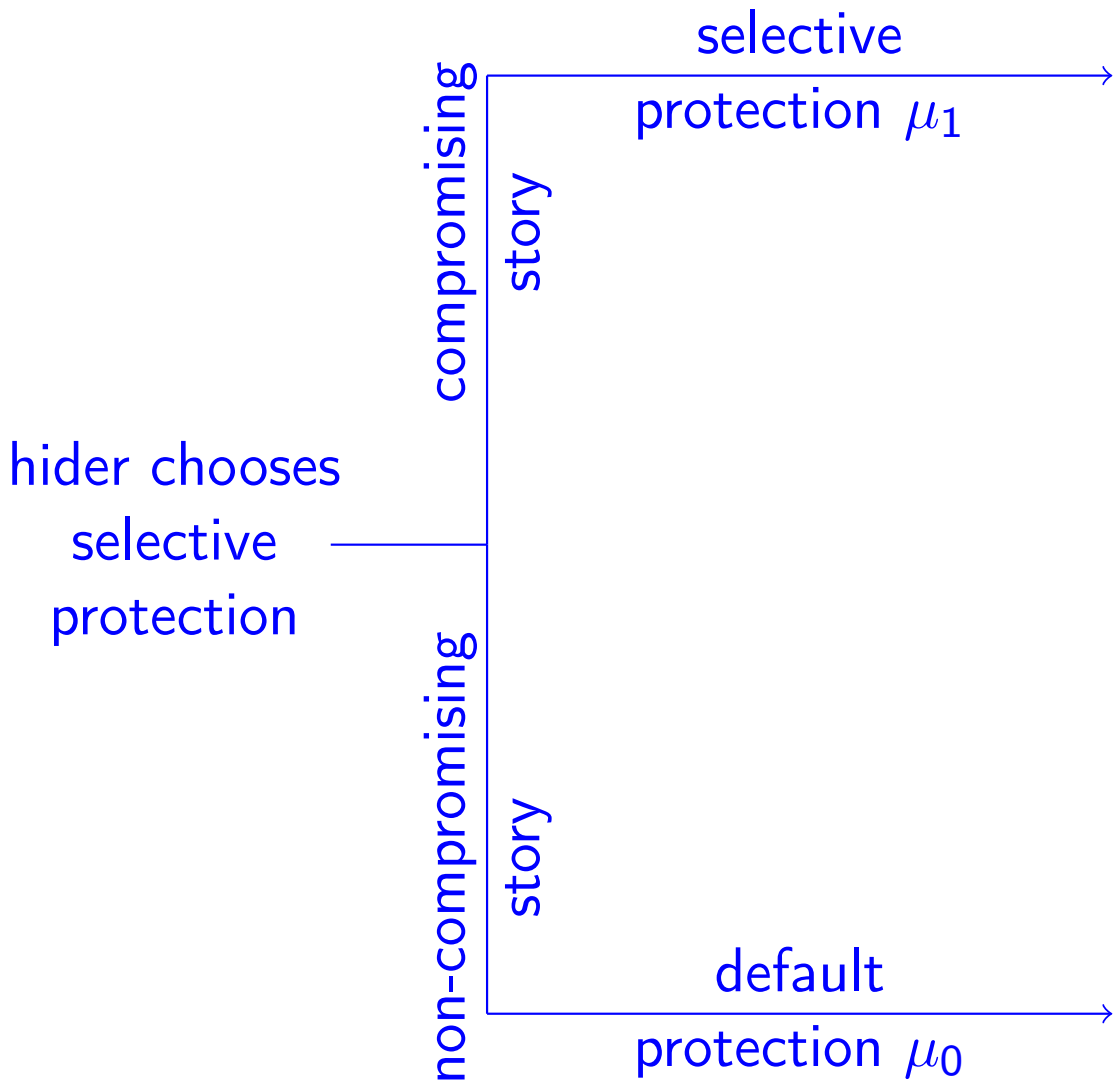
selective
protection μ_1
compromising
story
hider chooses
selective
protection
non-compromising
story
default
protection μ_0

Model

- Hider commits to the level of selective protection, parametrized by $\mu_1 > 0$. The level of default protection $\mu_0 > 0$ is exogenous.
- Hider gets involved in a story of type $\theta \in \{0, 1\}$
  - $\theta = 1$: compromising story
  - $\theta = 0$: non-compromising story

  Seeker knows μ_0 and μ_1, and believes $\Pr(\theta = 1) = p$
- Seeker learns the story through Poisson process with rate μ_θ and at flow cost $c > 0$.

Model

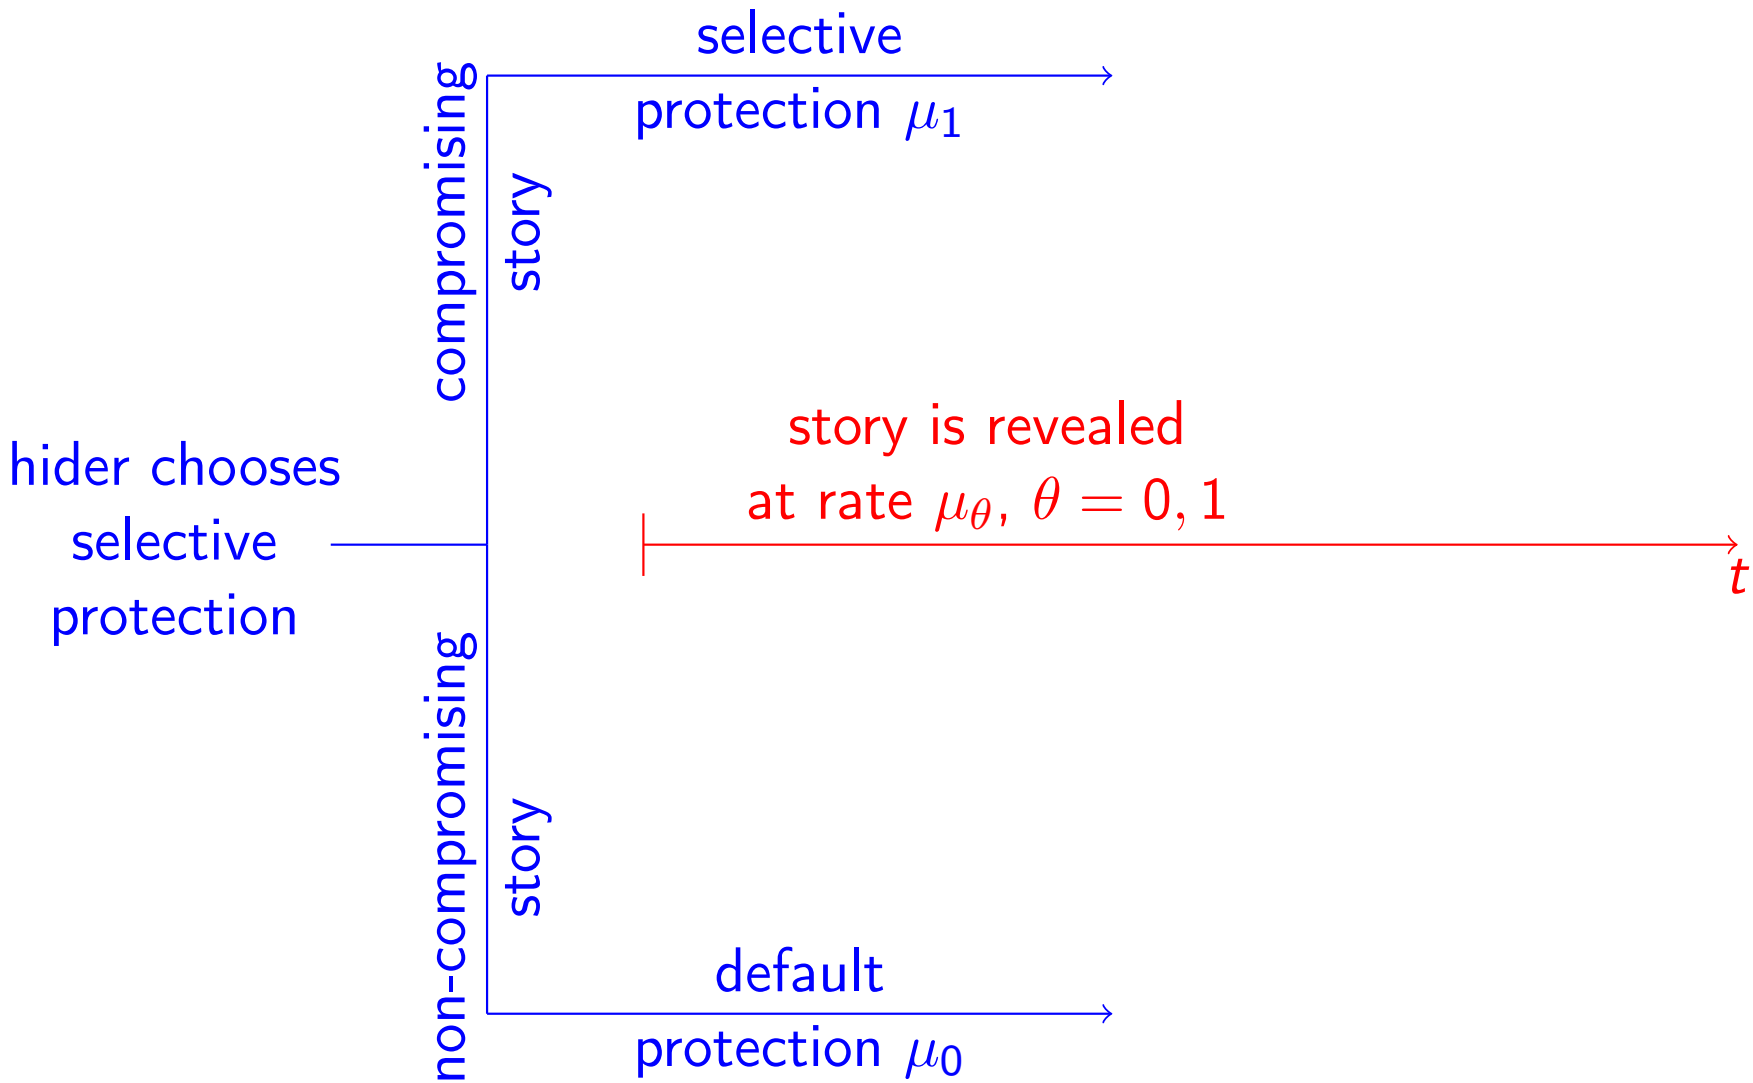
selective
protection μ_1
compromising
story
hider chooses
selective
protection
story is revealed
at rate μ_θ, $\theta = 0, 1$
t
non-compromising
story
default
protection μ_0

Model

- Hider commits to the level of selective protection, parametrized by $\mu_1 > 0$. The level of default protection $\mu_0 > 0$ is exogenous.
- Hider gets involved in a story of type $\theta \in \{0, 1\}$
  - $\theta = 1$: compromising story
  - $\theta = 0$: non-compromising story

  Seeker knows μ_0 and μ_1, and believes $\Pr(\theta = 1) = p$
- Seeker learns the story through Poisson process with rate μ_θ and at flow cost $c > 0$.
- Seeker gets 1 if he reports a compromising story and a negative payoff if he reports a non-compromising story. To report the story, the seeker has to learn it.

Model

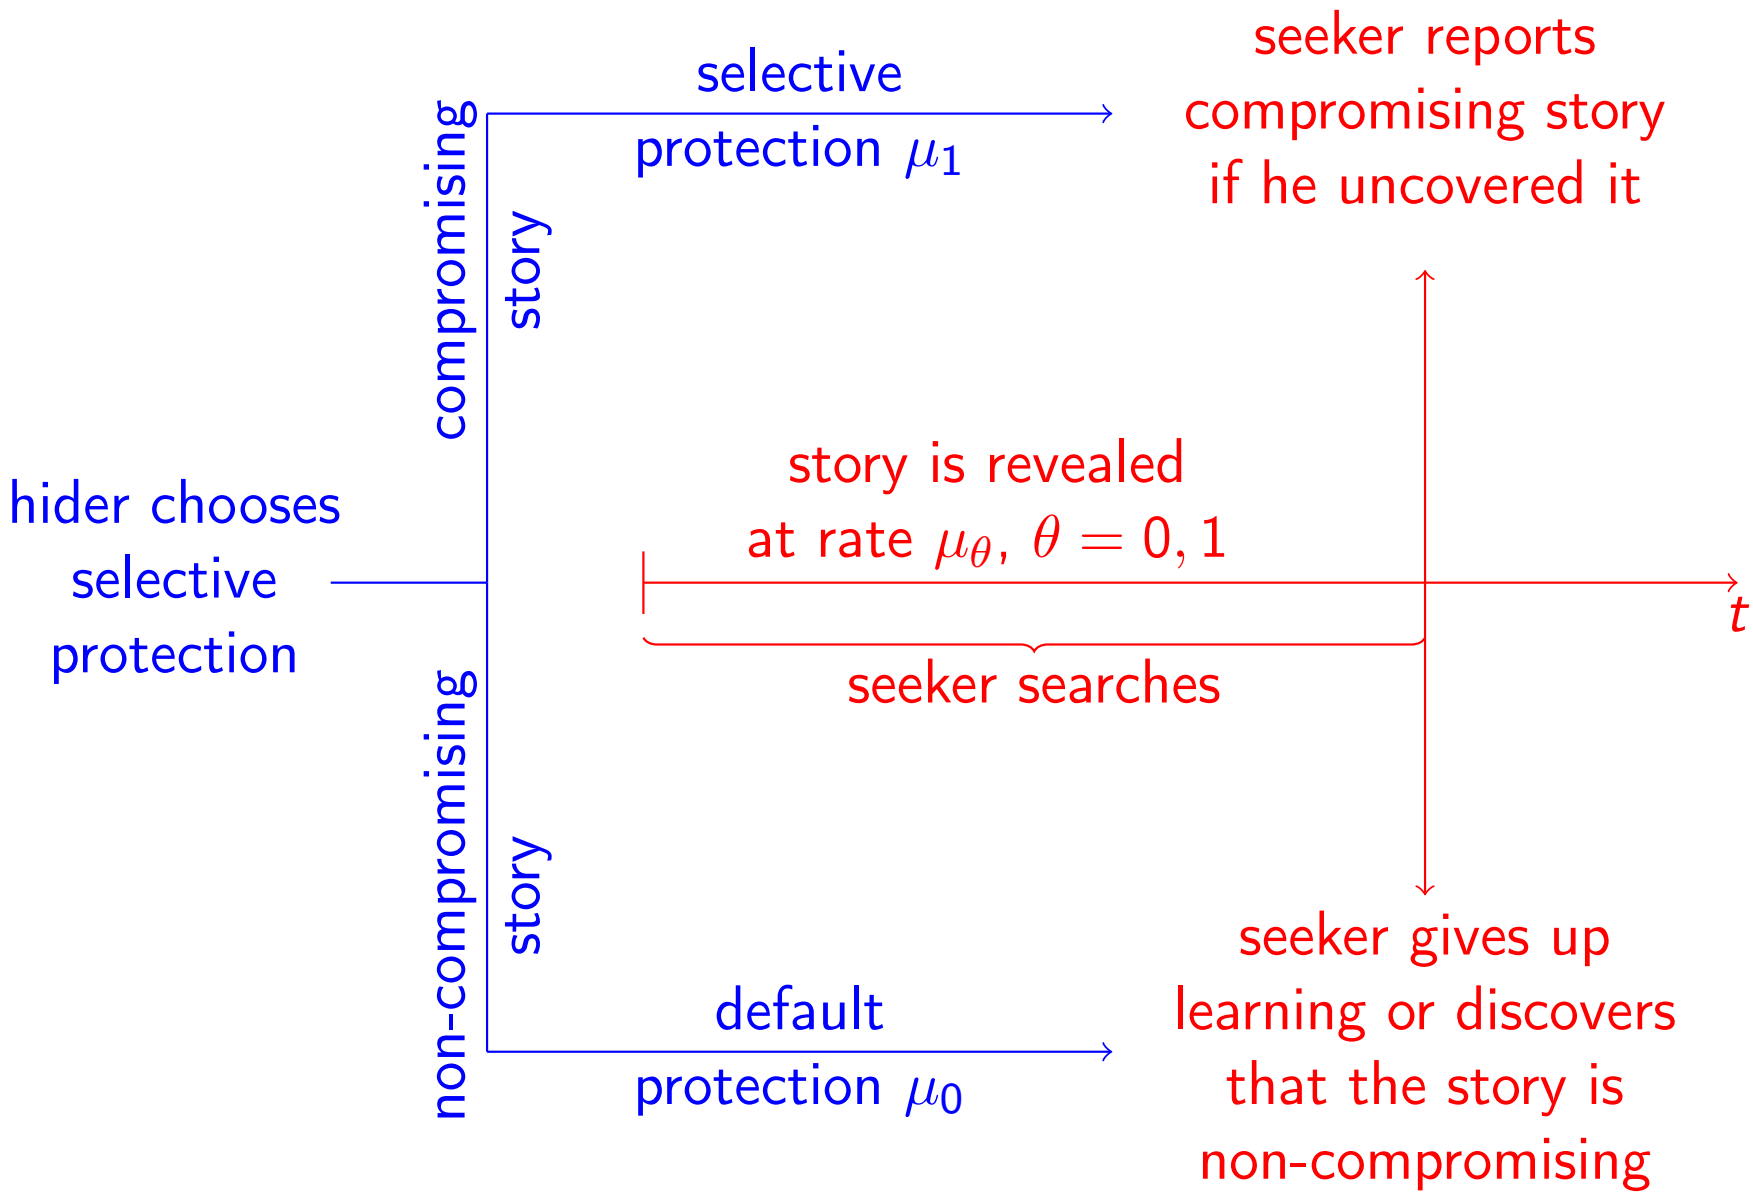
hider chooses
selective
protection
compromising
story
selective
protection μ_1
seeker reports
compromising story
if he uncovered it
story is revealed
at rate μ_θ, $\theta = 0, 1$
t
seeker searches
non-compromising
story
default
protection μ_0
seeker gives up
learning or discovers
that the story is
non-compromising

Model

- Hider commits to the level of selective protection, parametrized by $\mu_1 > 0$. The level of default protection $\mu_0 > 0$ is exogenous.
- Hider gets involved in a story of type $\theta \in \{0, 1\}$
  - $\theta = 1$: compromising story
  - $\theta = 0$: non-compromising story

  Seeker knows μ_0 and μ_1, and believes $\Pr(\theta = 1) = p$
- Seeker learns the story through Poisson process with rate μ_θ and at flow cost $c > 0$.
- Seeker gets 1 if he reports a compromising story and a negative payoff if he reports a non-compromising story. To report the story, the seeker has to learn it.
- Hider maximizes Pr the seeker fails to report the story.
  - $\Rightarrow$ cost of protection $= 0$

Model

- Hider commits to the level of selective protection, parametrized by $\mu_1 > 0$. The level of default protection $\mu_0 > 0$ is exogenous.
- Hider gets involved in a story of type $\theta \in \{0, 1\}$
  - $\theta = 1$: compromising story
  - $\theta = 0$: non-compromising story

  Seeker knows μ_0 and μ_1, and believes $\Pr(\theta = 1) = p$
- Seeker learns the story through Poisson process with rate μ_θ and at flow cost $c > 0$.
- Seeker gets 1 if he reports a compromising story and a negative payoff if he reports a non-compromising story. To report the story, the seeker has to learn it.
- Hider maximizes Pr the seeker fails to report the story.
  - $\Rightarrow$ cost of protection $= 0$
- No discounting

Model

- Hider commits to the level of selective protection, parametrized by $\mu_1 > 0$. The level of default protection $\mu_0 > 0$ is exogenous.
- Hider gets involved in a story of type $\theta \in \{0, 1\}$
  - $\theta = 1$: compromising story
  - $\theta = 0$: non-compromising story

  Seeker knows μ_0 and μ_1, and believes $\Pr(\theta = 1) = p$
- Seeker learns the story through Poisson process with rate μ_θ and at flow cost $c > 0$.
- Seeker gets 1 if he reports a compromising story and a negative payoff if he reports a non-compromising story. To report the story, the seeker has to learn it.
- Hider maximizes Pr the seeker fails to report the story.
  - $\Rightarrow$ cost of protection $= 0$
- No discounting

Assumption

Hider can choose any μ_1 such that $\mu_1 \geq c/p$

Results

If $\mu_1 < \mu_0$:

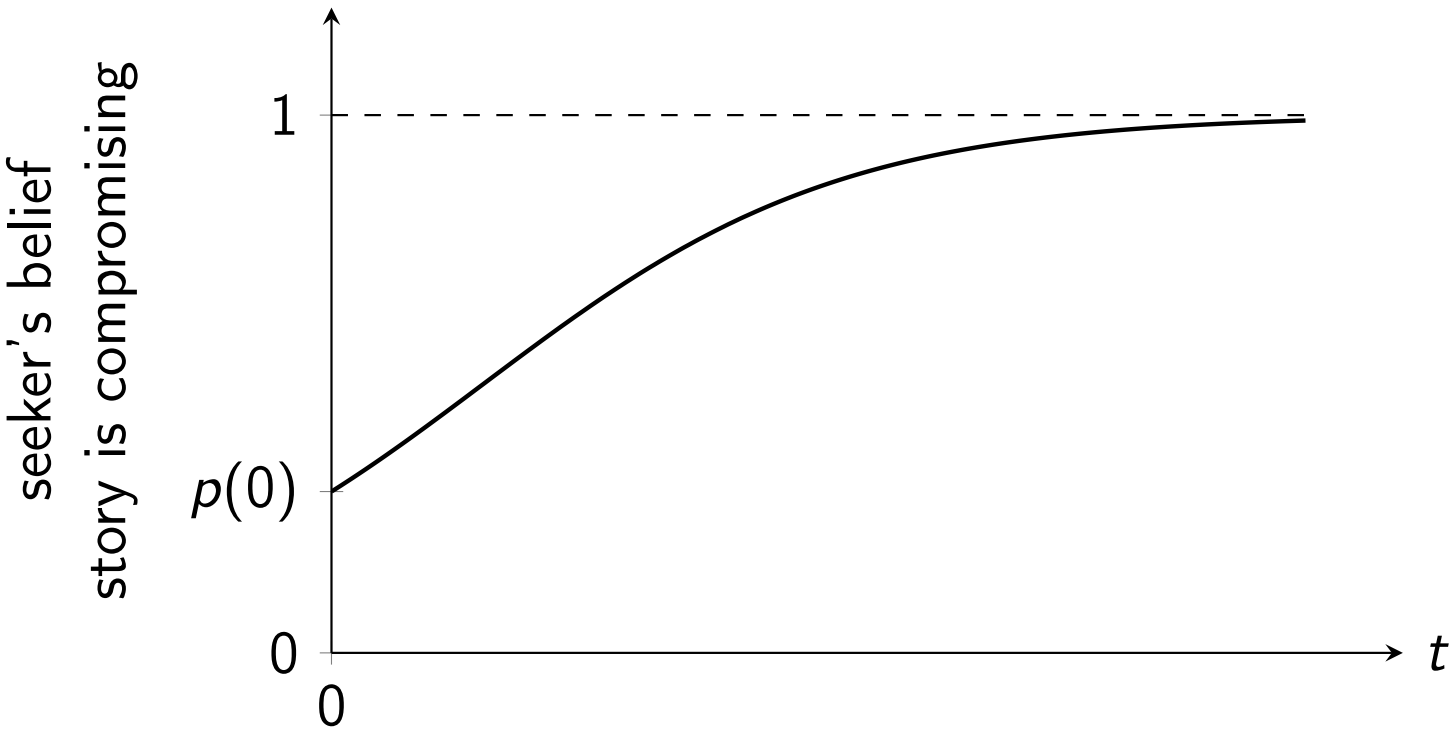

Results

If $\mu_1 < \mu_0$:

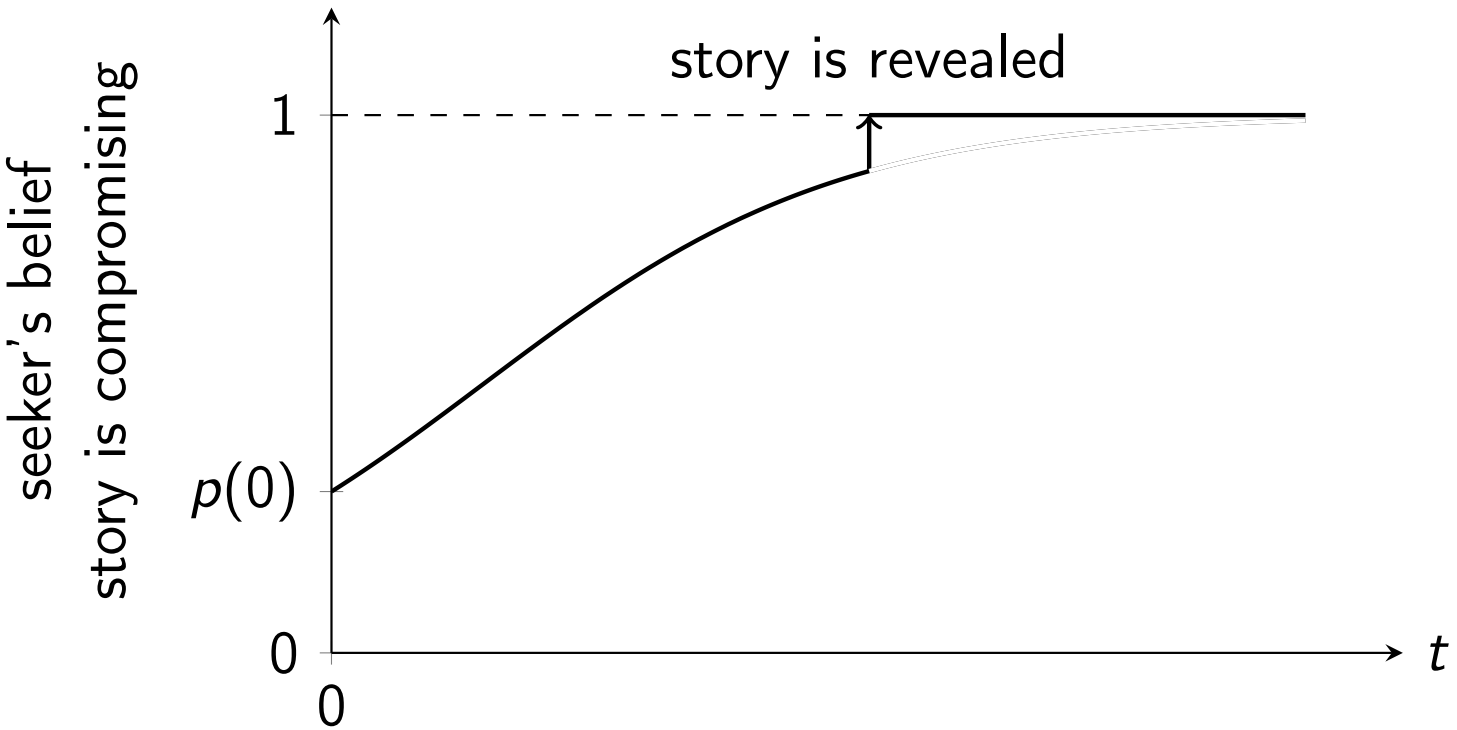

Results

If $c/p \leq \mu_1 < \mu_0$: learning is optimal at $p(0)$

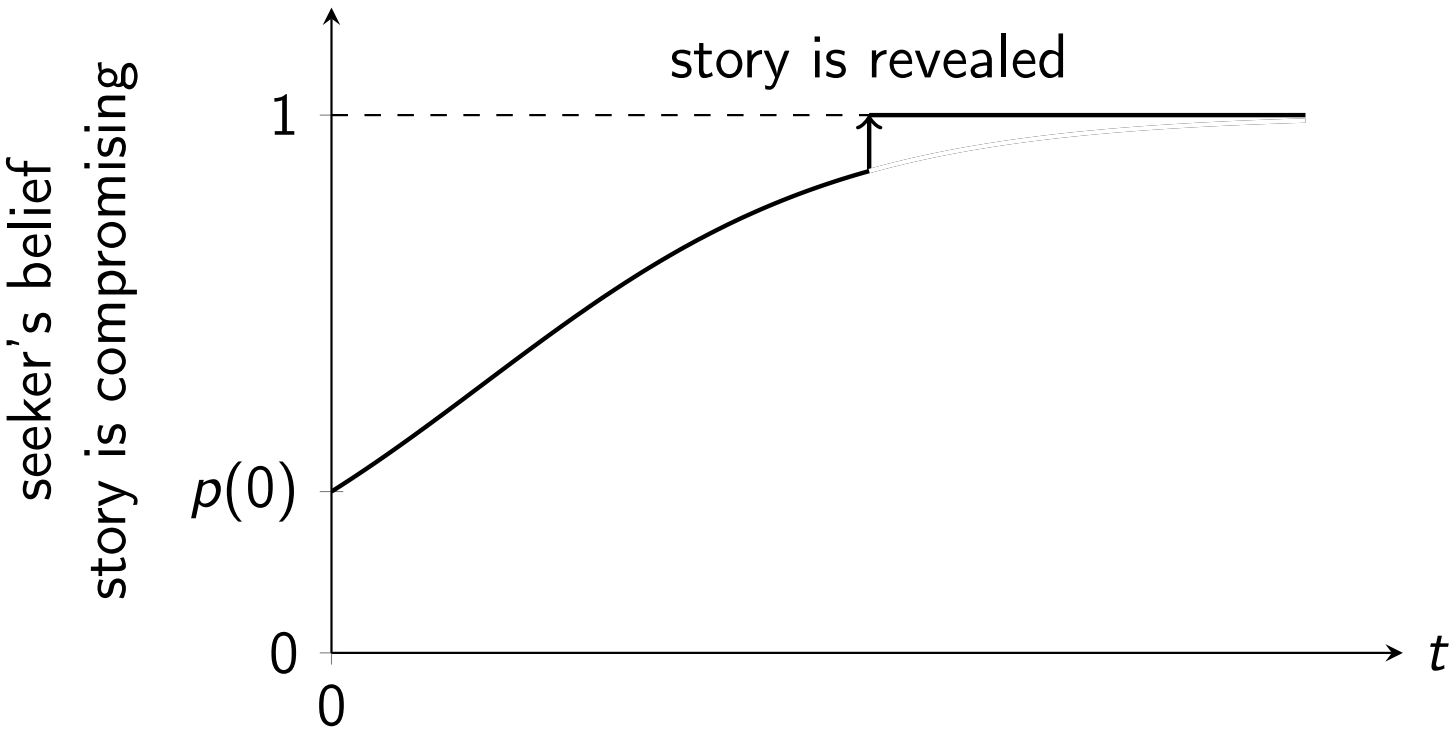

If $\mu_1 > \mu_0 \geq c/p$: seeker stops learning in finite time

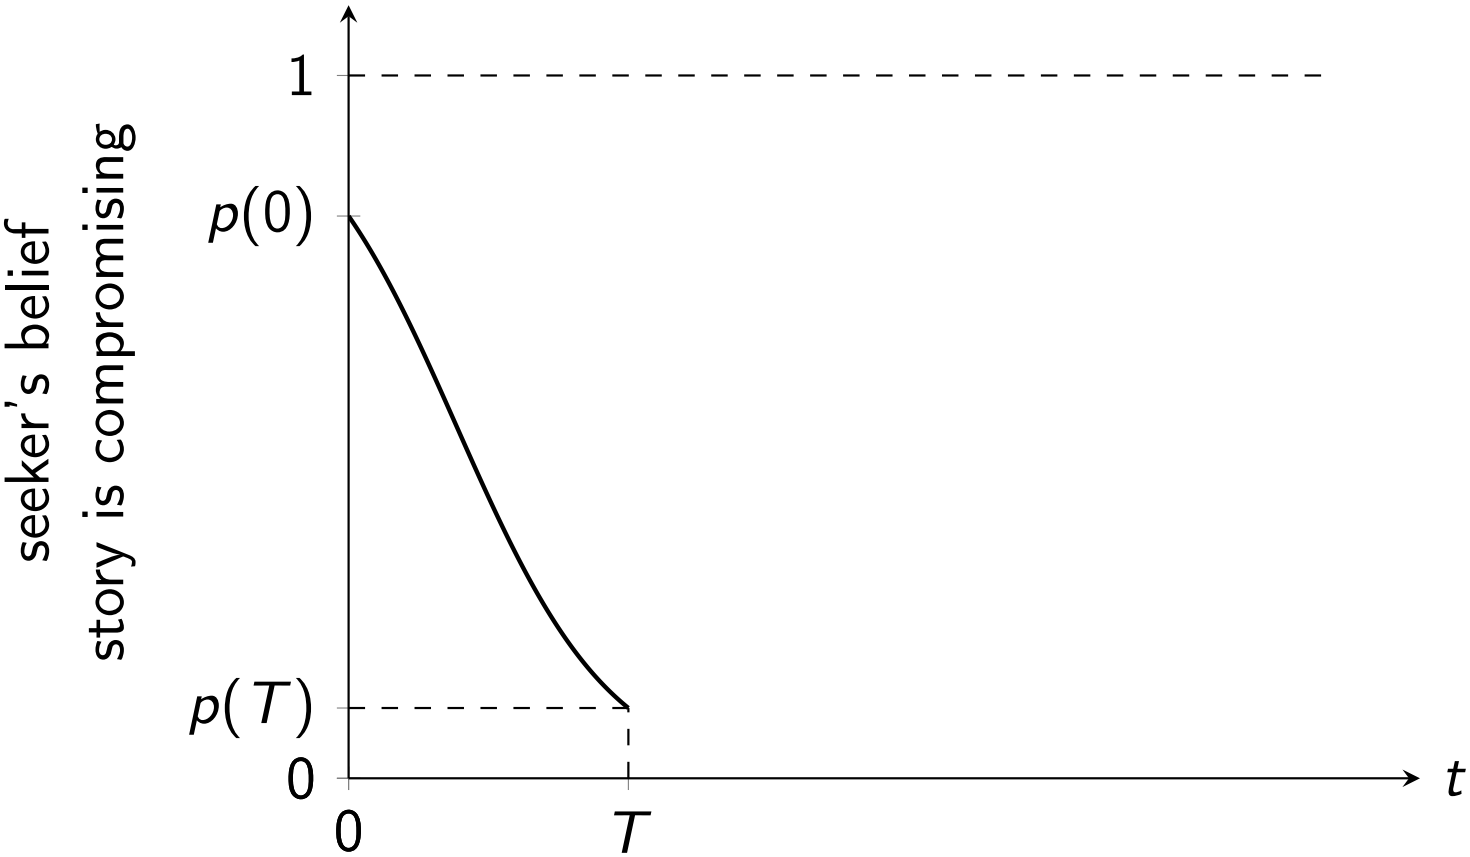

If $\mu_1 > \mu_0 \geq c/p$: seeker stops learning in finite time

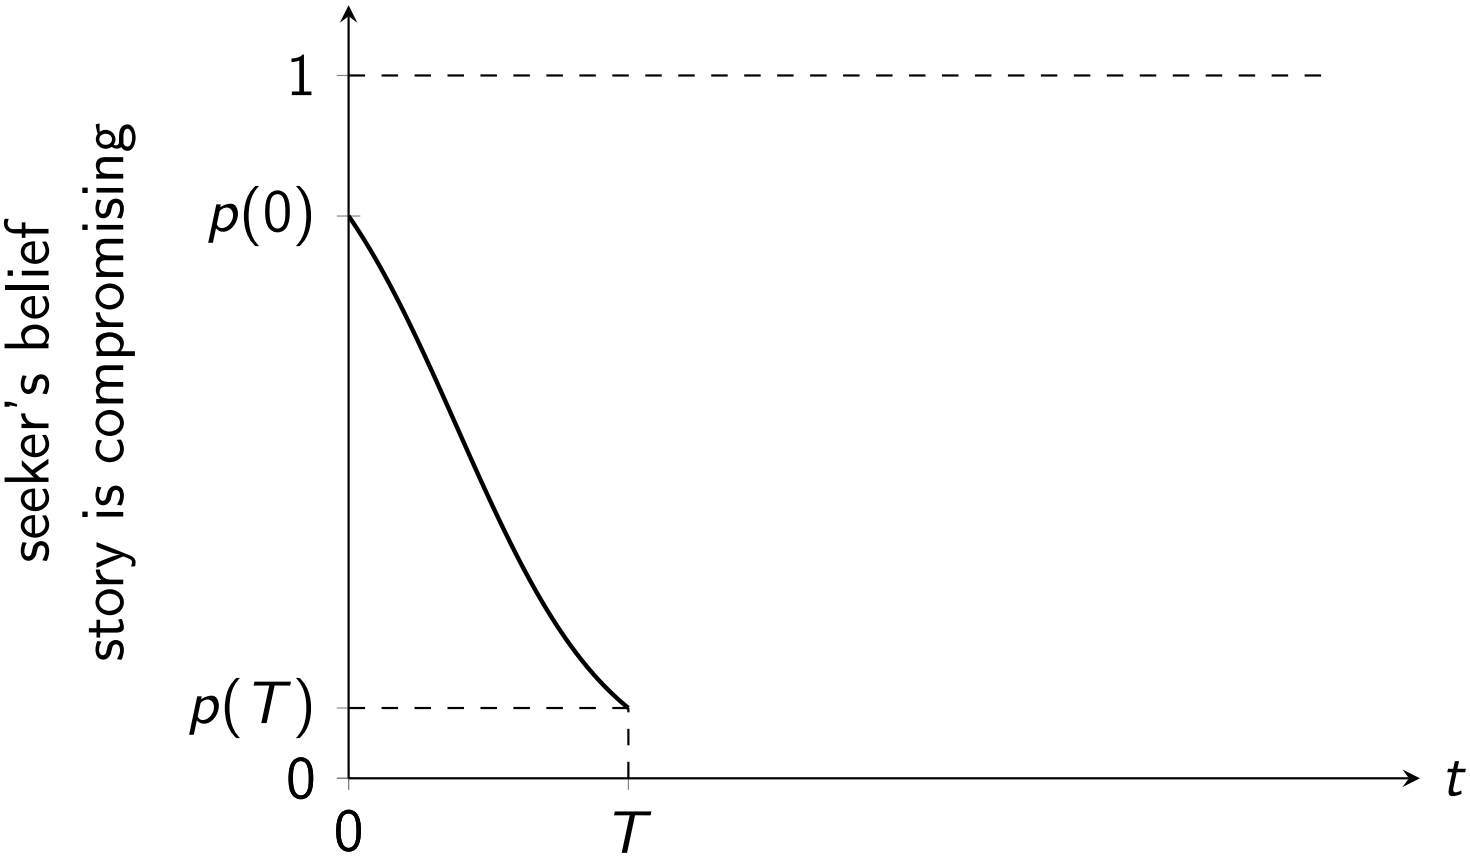

Result 1

If $\mu_0 > c/p$ (default protection is weak), then the optimal selective protection is even weaker: $\mu_1 > \mu_0$ at the optimum.

Trade-off

Trade-off

Trade-off

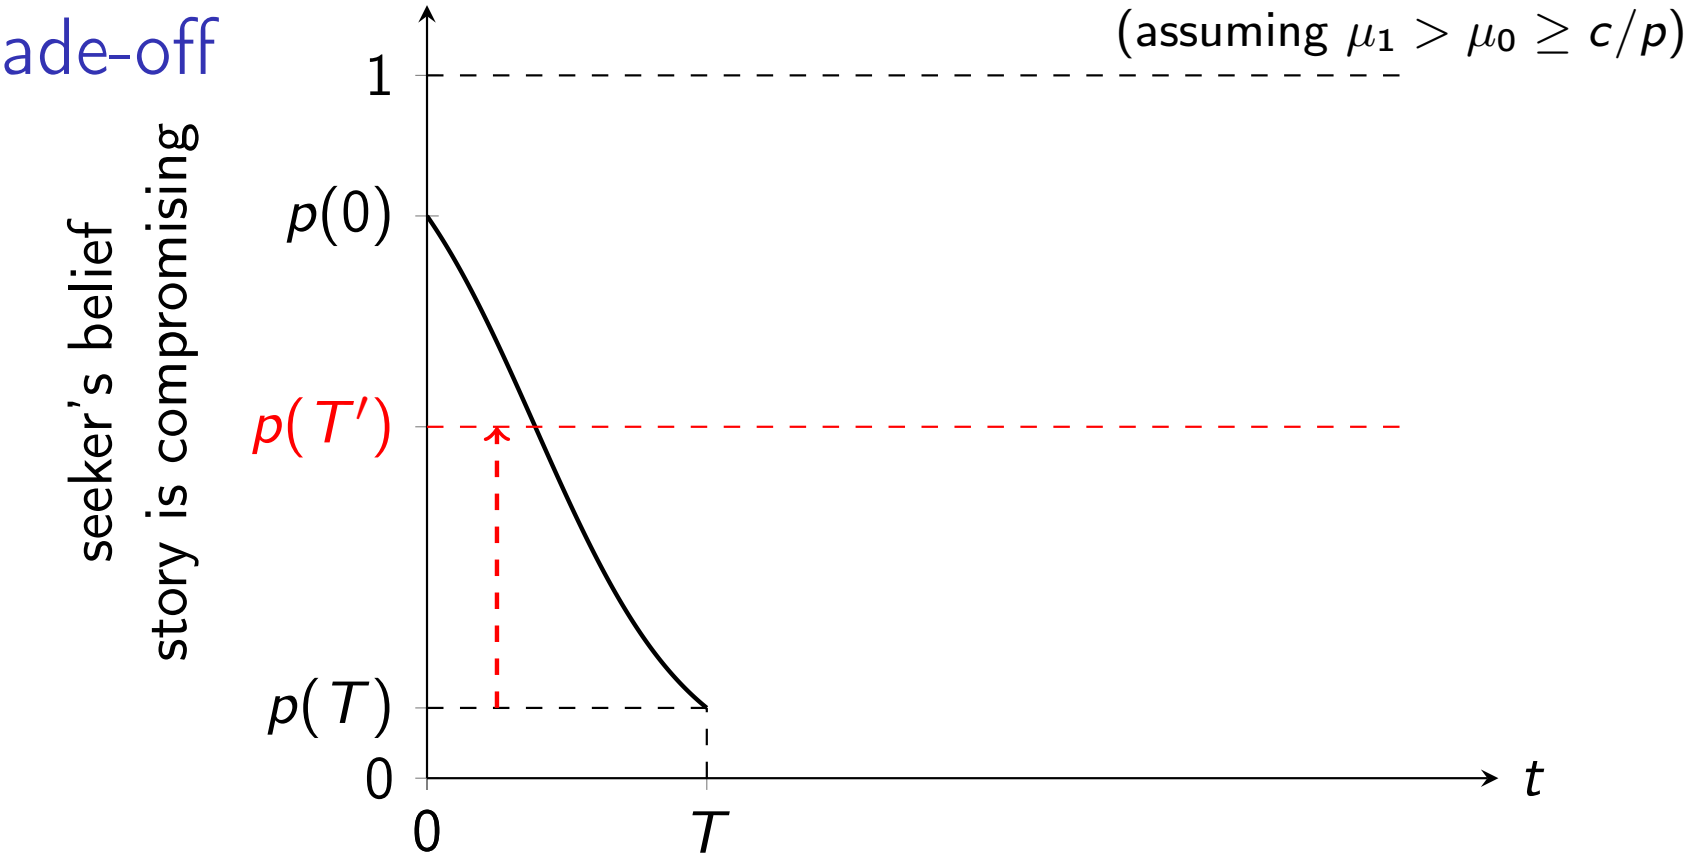
(assuming $\mu_1 > \mu_0 \geq c/p$)
1
$p(0)$
$p(T')$
$p(T)$
0
0
T
t
seeker's belief
story is compromising


Pr(no report)
cost effect (+)
selective protection

Trade-off

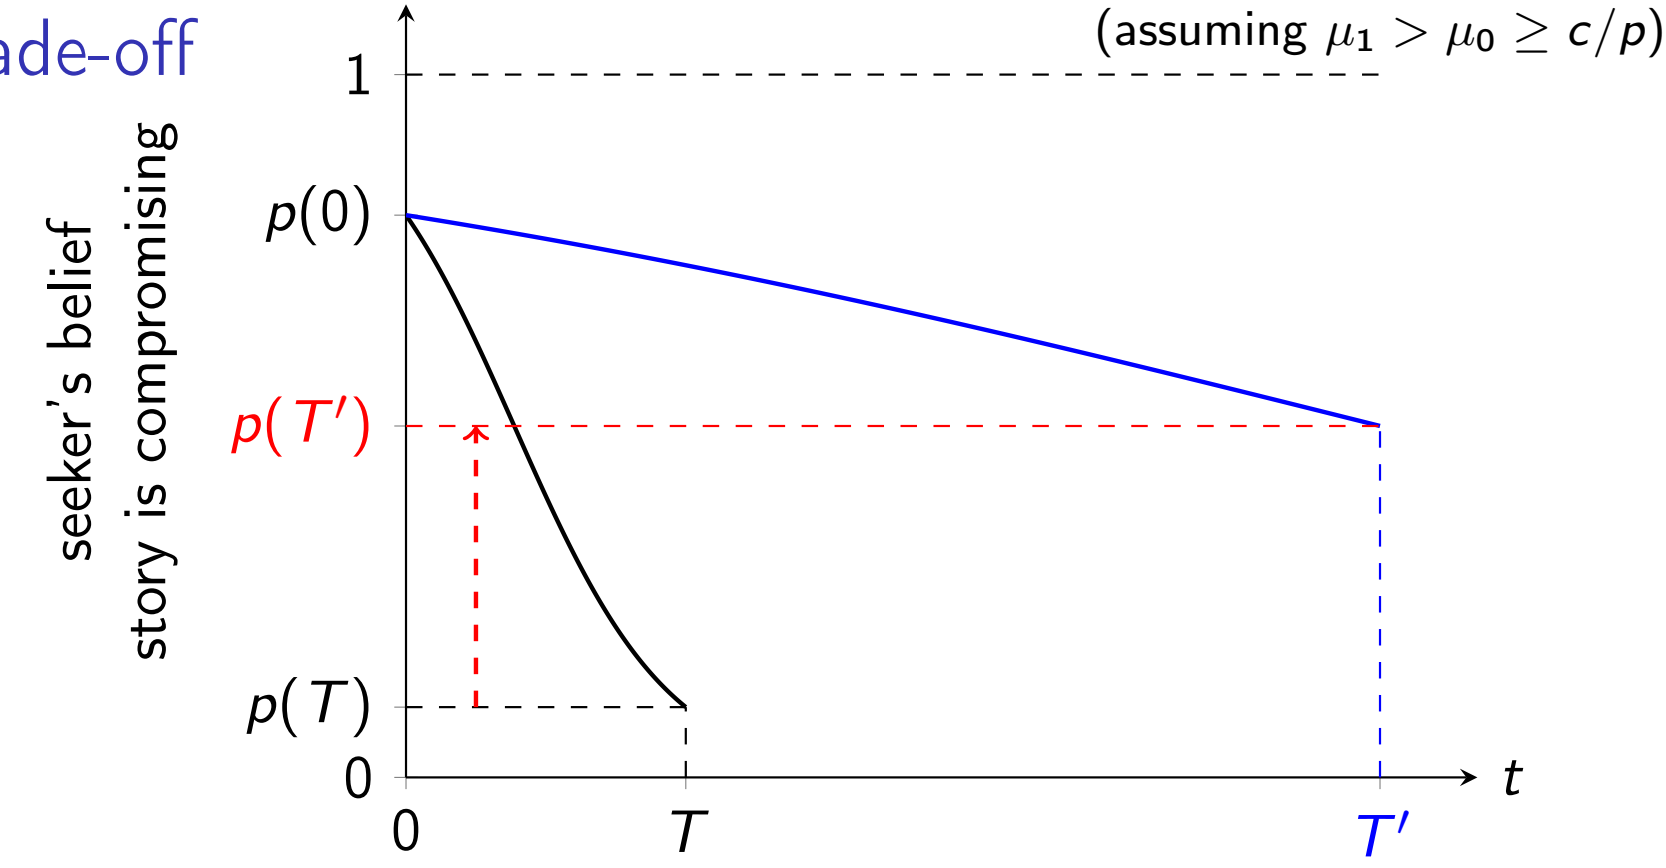
(assuming $\mu_1 > \mu_0 \geq c/p$)
1
$p(0)$
$p(T')$
$p(T)$
0
seeker's belief
story is compromising
t
0
T
T'

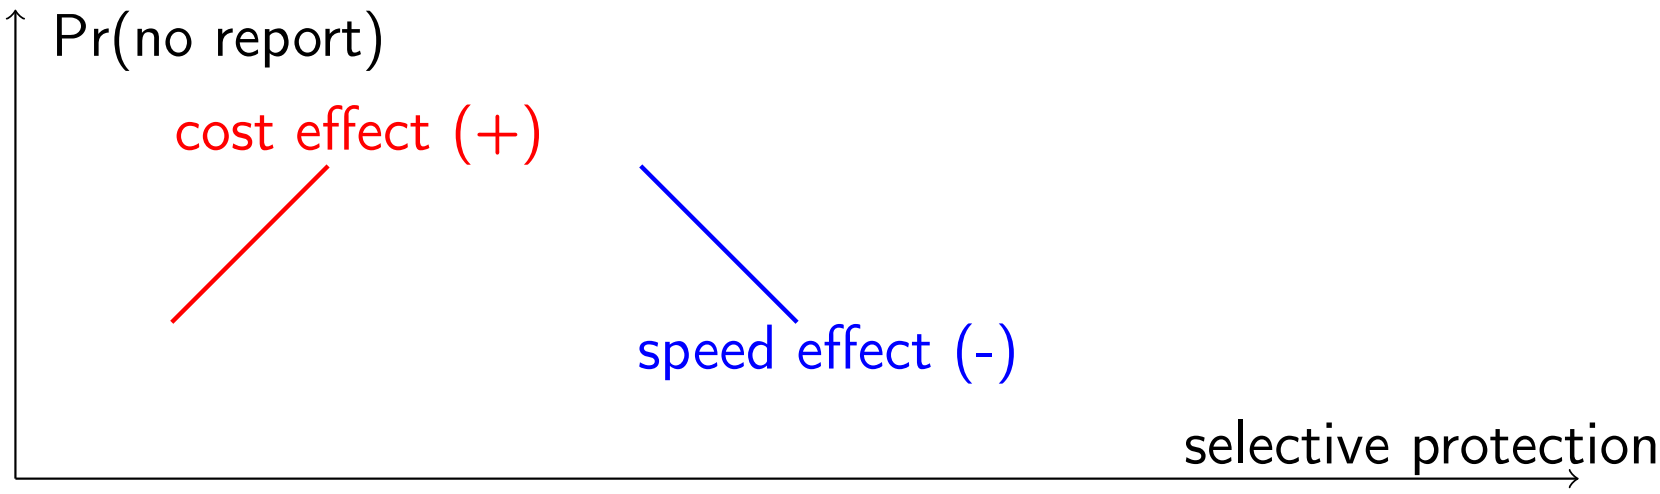
Pr(no report)
cost effect (+)
speed effect (-)
selective protection

Trade-off

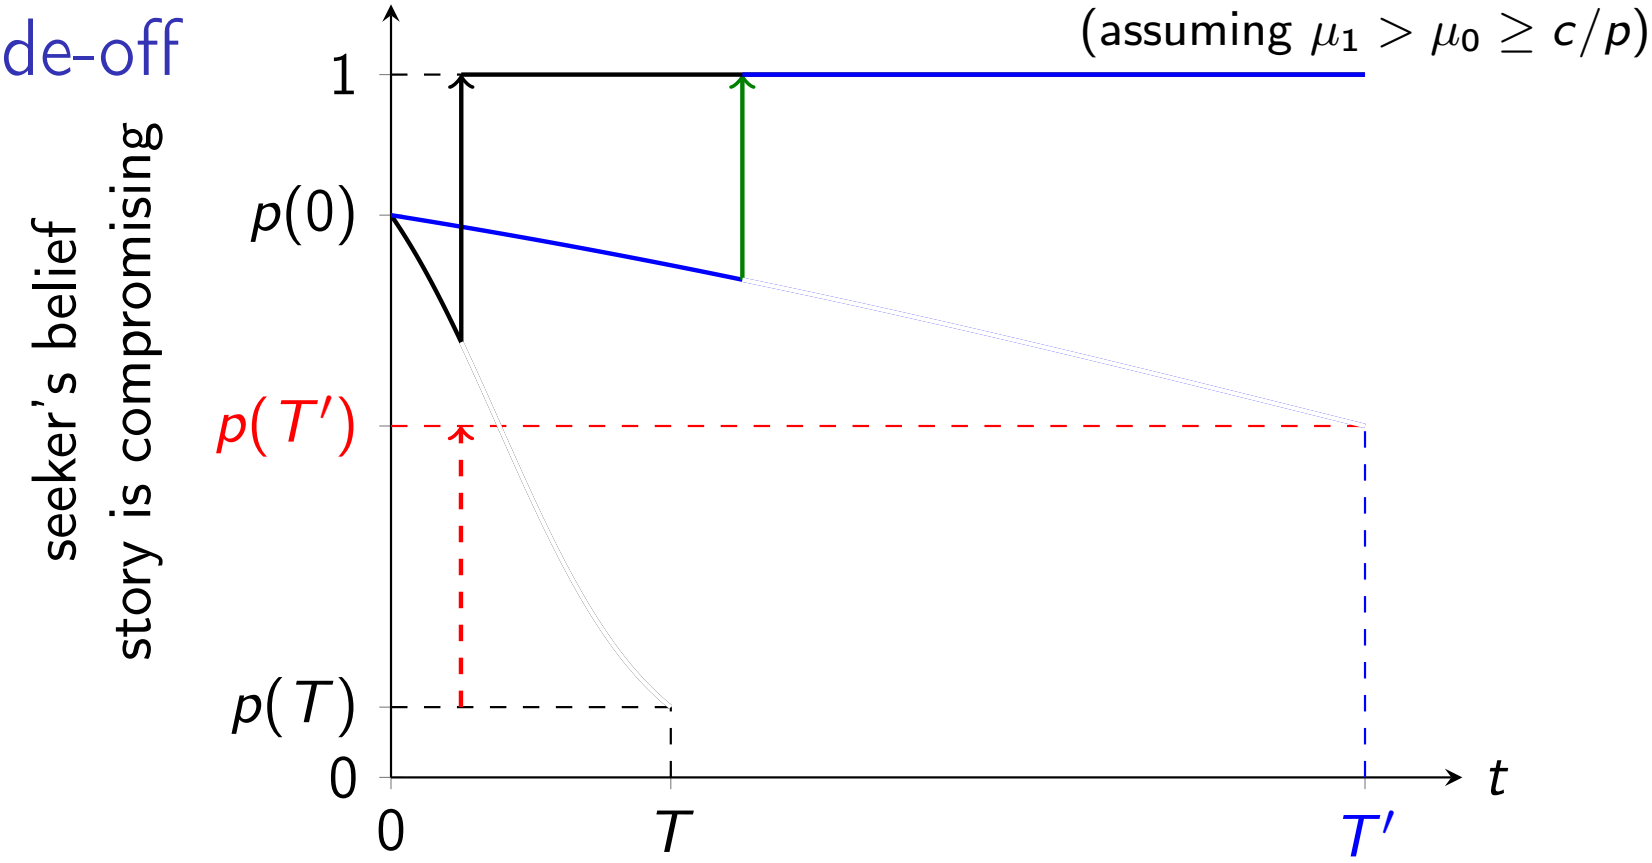
(assuming $\mu_1 > \mu_0 \geq c/p$)
seeker's belief
story is compromising
1
$p(0)$
$p(T')$
$p(T)$
0
0
T
T'
t

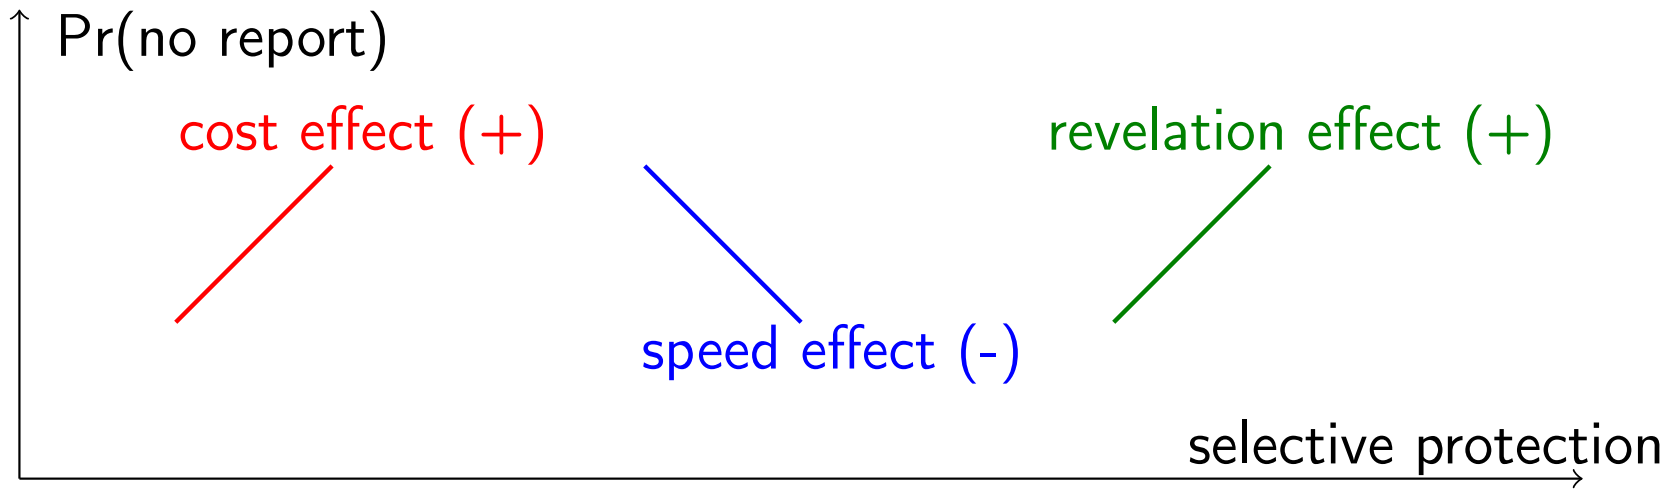
Pr(no report)
cost effect (+)
revelation effect (+)
speed effect (-)
selective protection

Trade-off

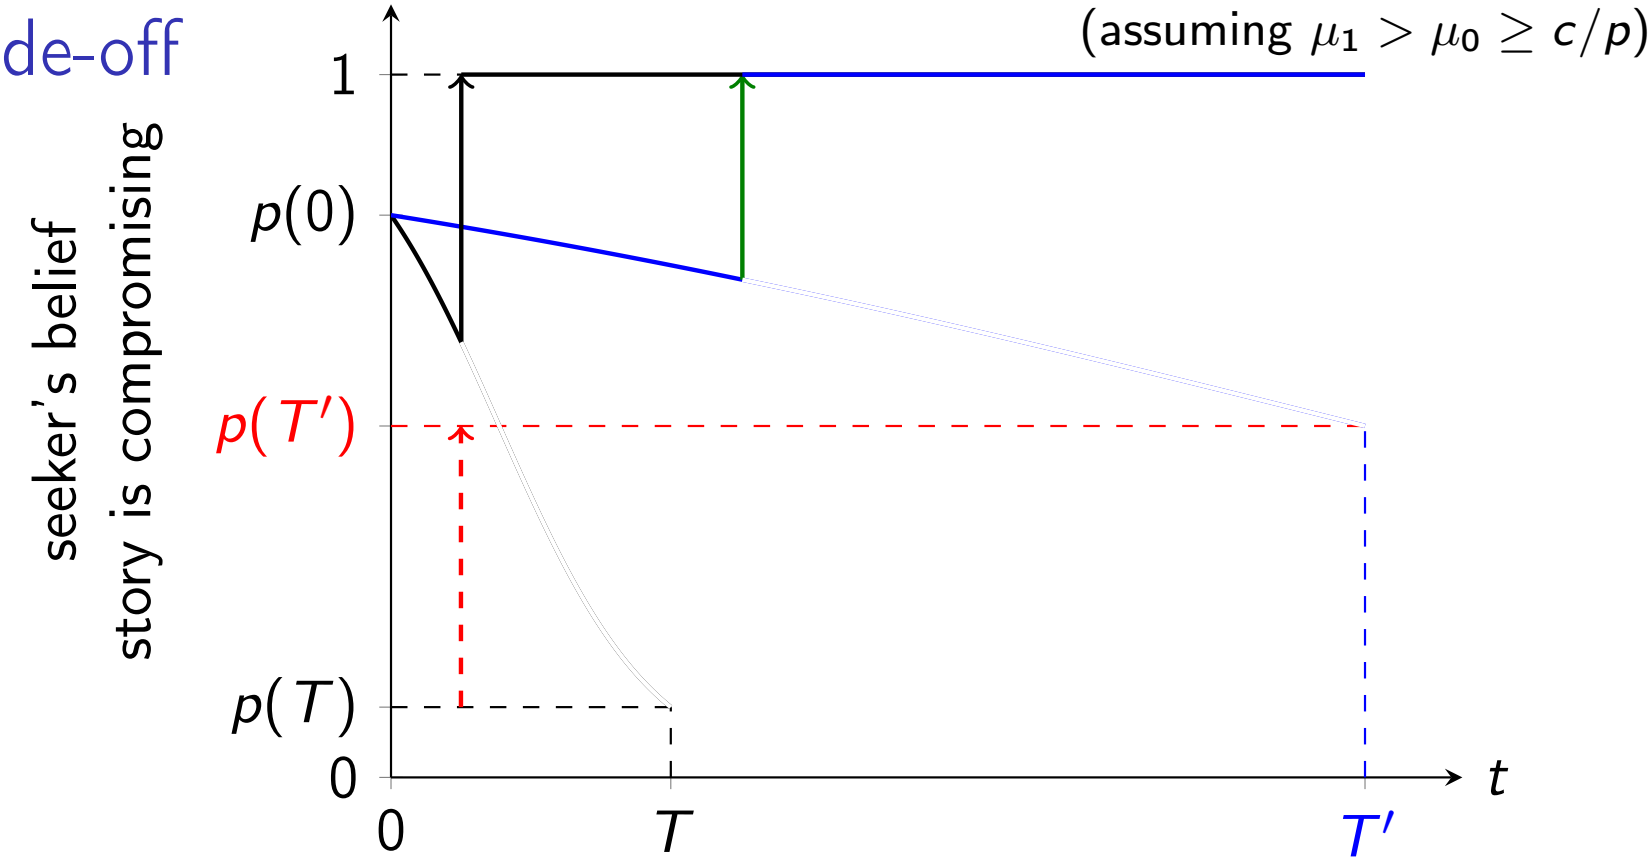
(assuming $\mu_1 > \mu_0 \geq c/p$)
1
$p(0)$
$p(T')$
$p(T)$
0
seeker's belief
story is compromising
t
0
T
T'

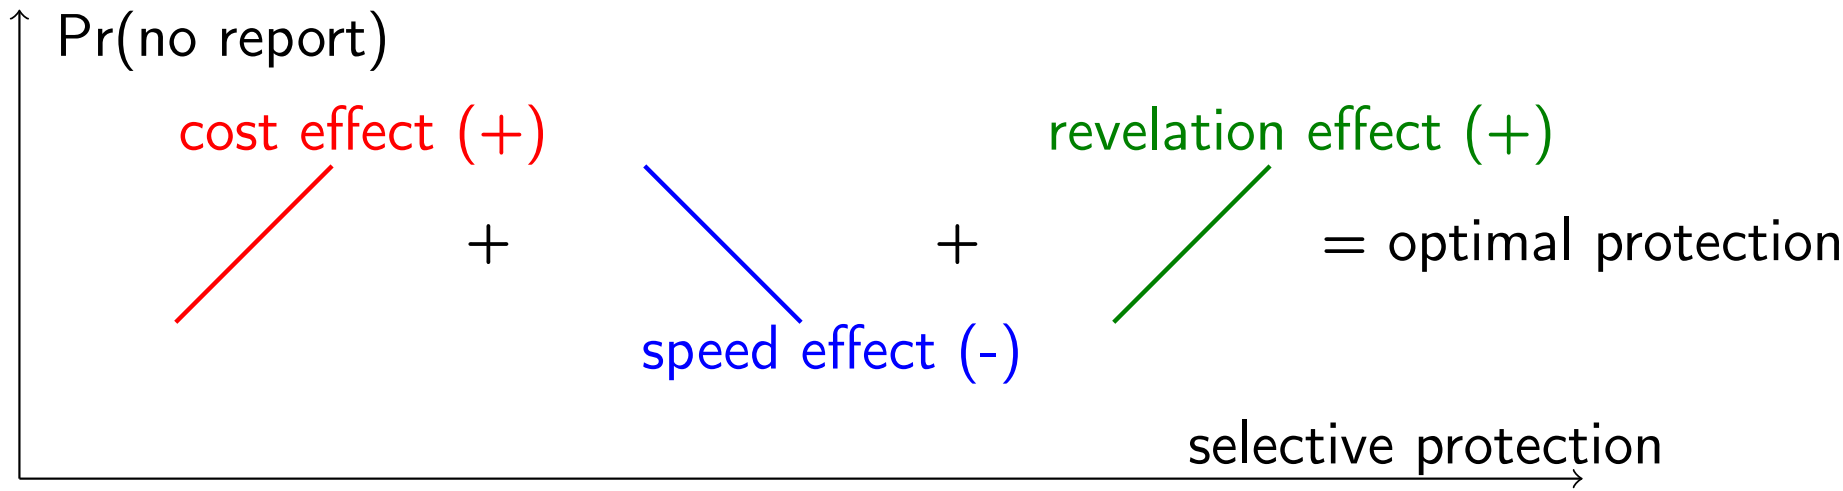
Pr(no report)
cost effect (+)
+
speed effect (-)
+
revelation effect (+)
= optimal protection
selective protection

Trade-off

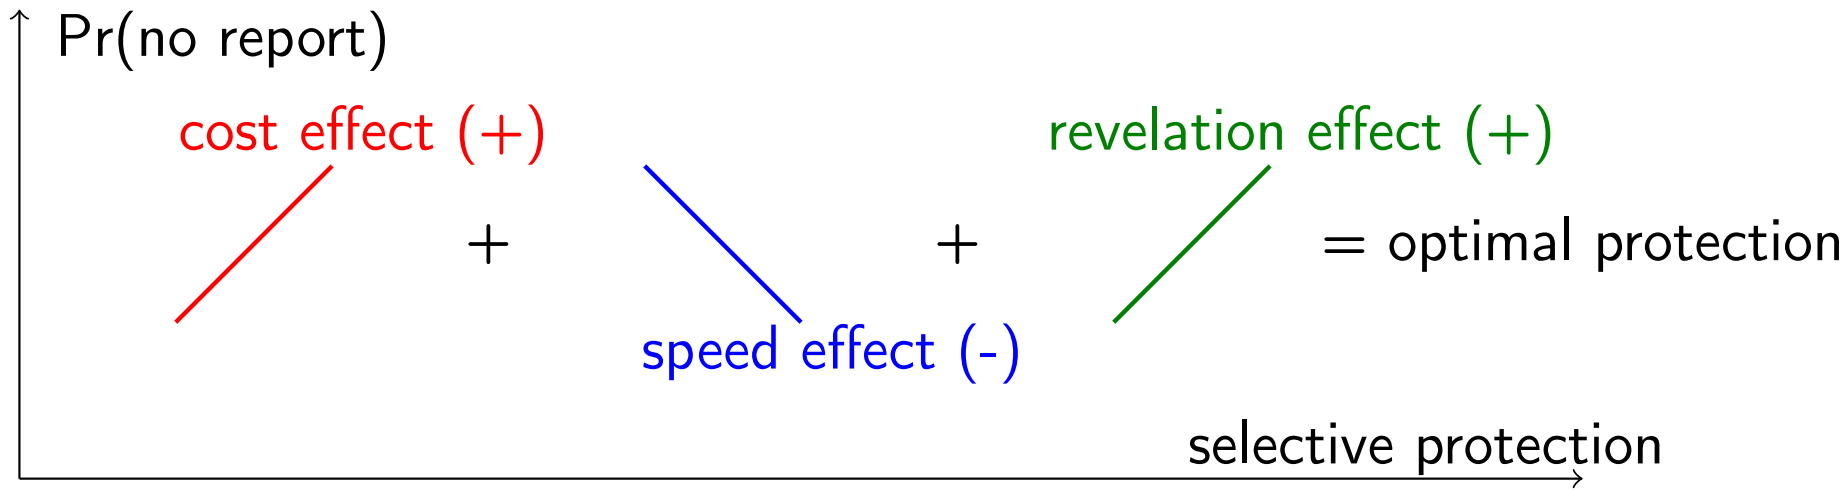
Pr(no report)
cost effect (+)
+
speed effect (-)
+
revelation effect (+)
= optimal protection
selective protection

Trade-off

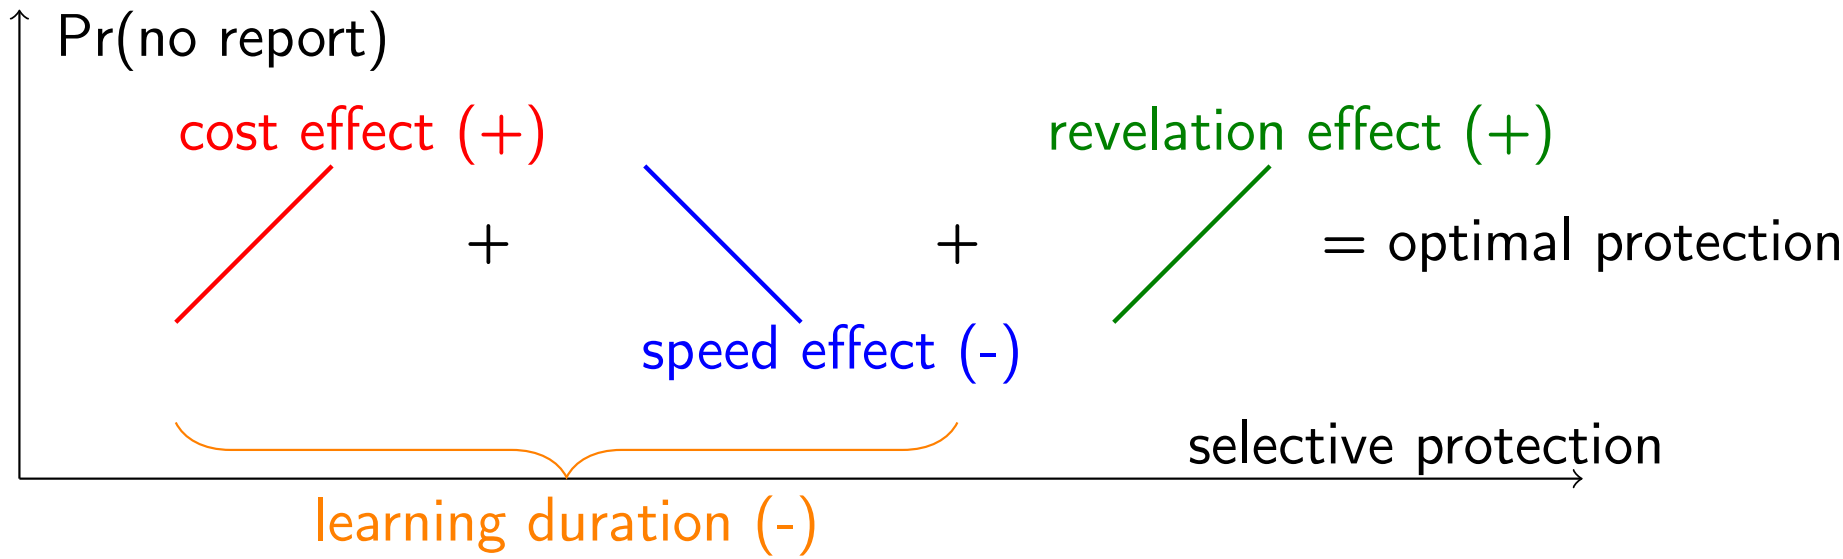
Pr(no report)
cost effect (+)
+
speed effect (-)
+
revelation effect (+)
= optimal protection
selective protection
learning duration (-)

Trade-off

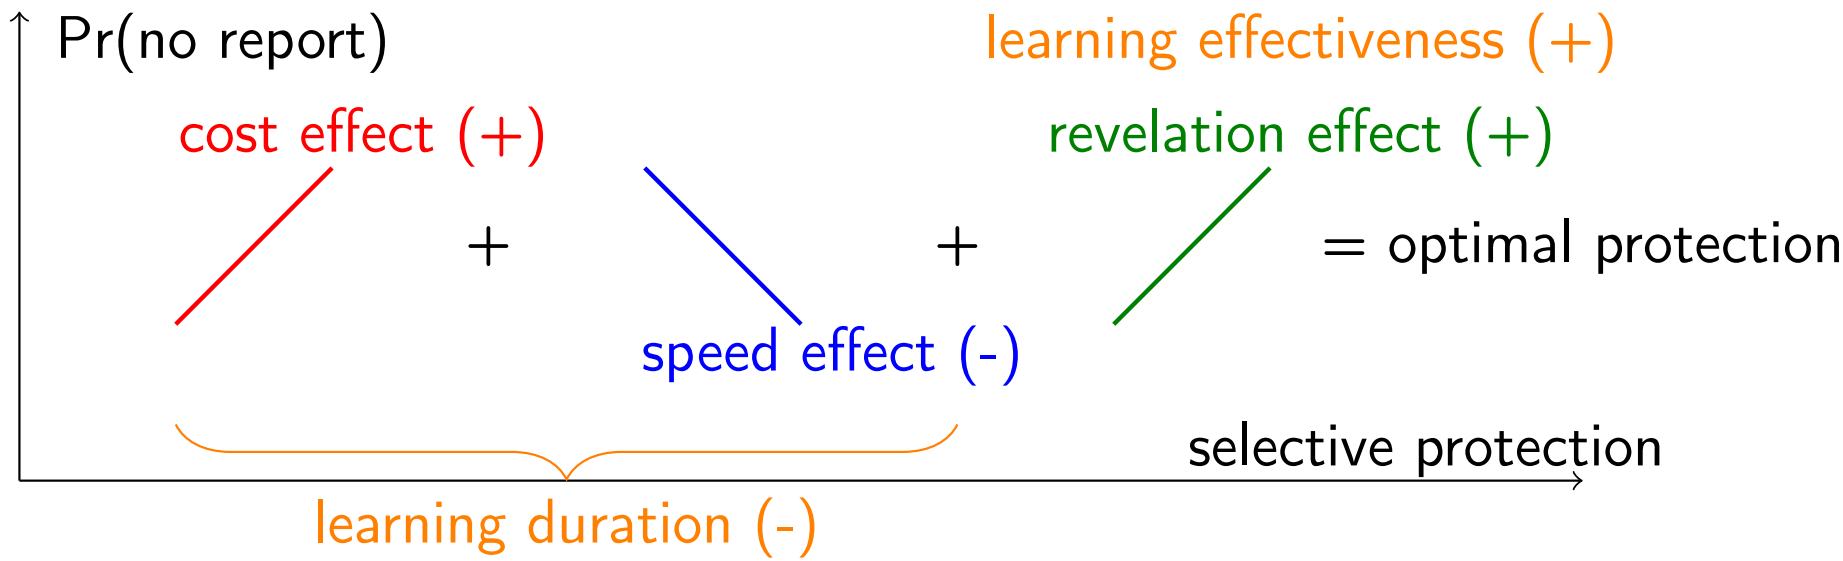
Pr(no report)
cost effect (+)
+
speed effect (-)
+
learning effectiveness (+)
revelation effect (+)
= optimal protection
selective protection
learning duration (-)

Optimal Protection

Theorem

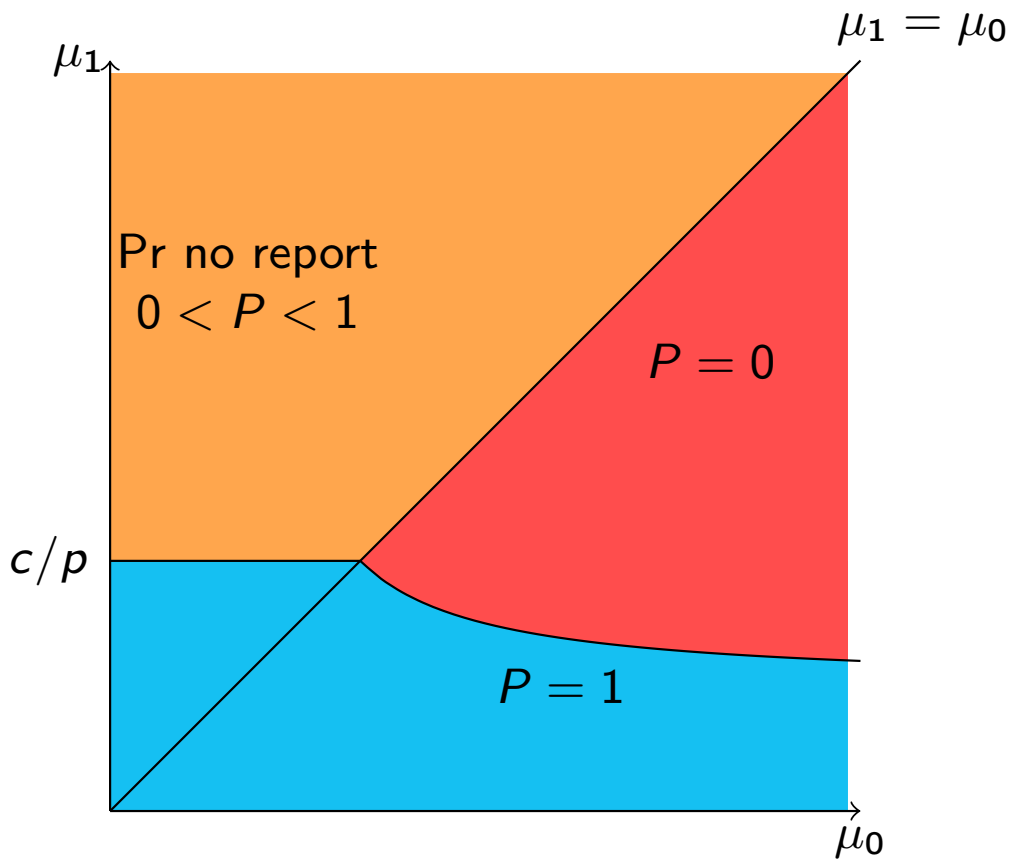

Optimal Protection

Theorem

Under assumption $\mu_1 \geq c/p$,

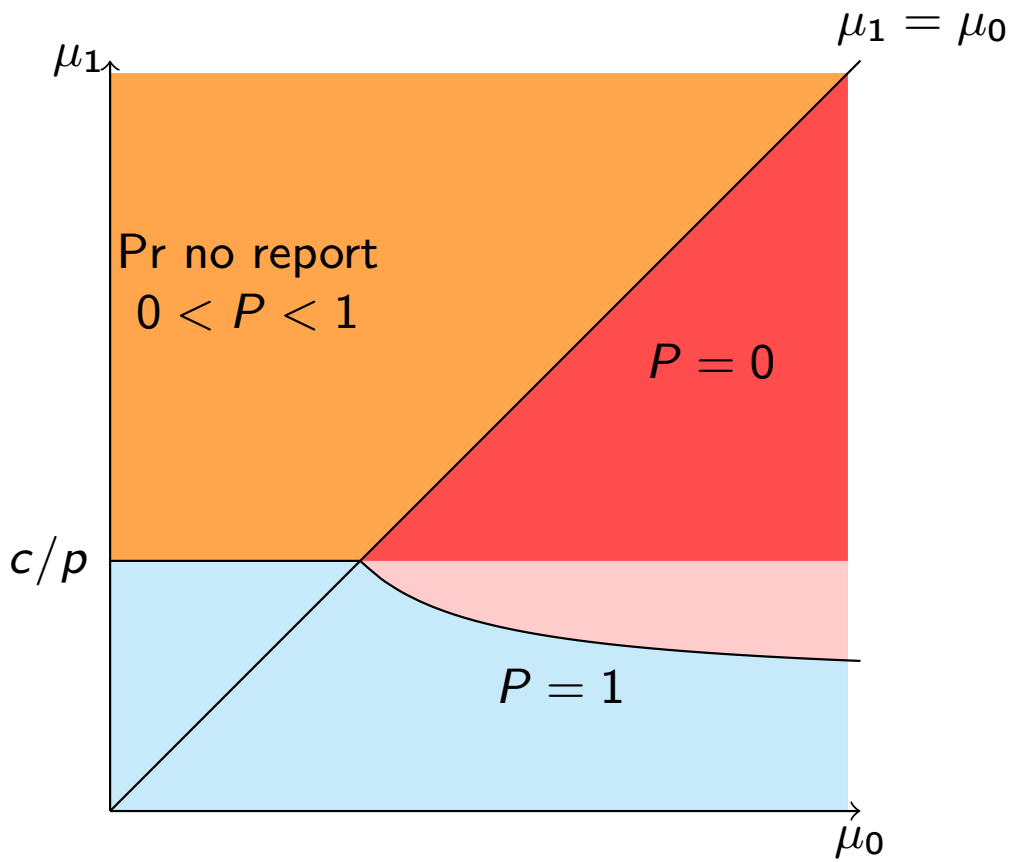

Optimal Protection

Theorem

Under assumption $\mu_1 \geq c/p$, the optimal selective protection is characterized by μ_1^*, which is a function of μ_0, c, and p.

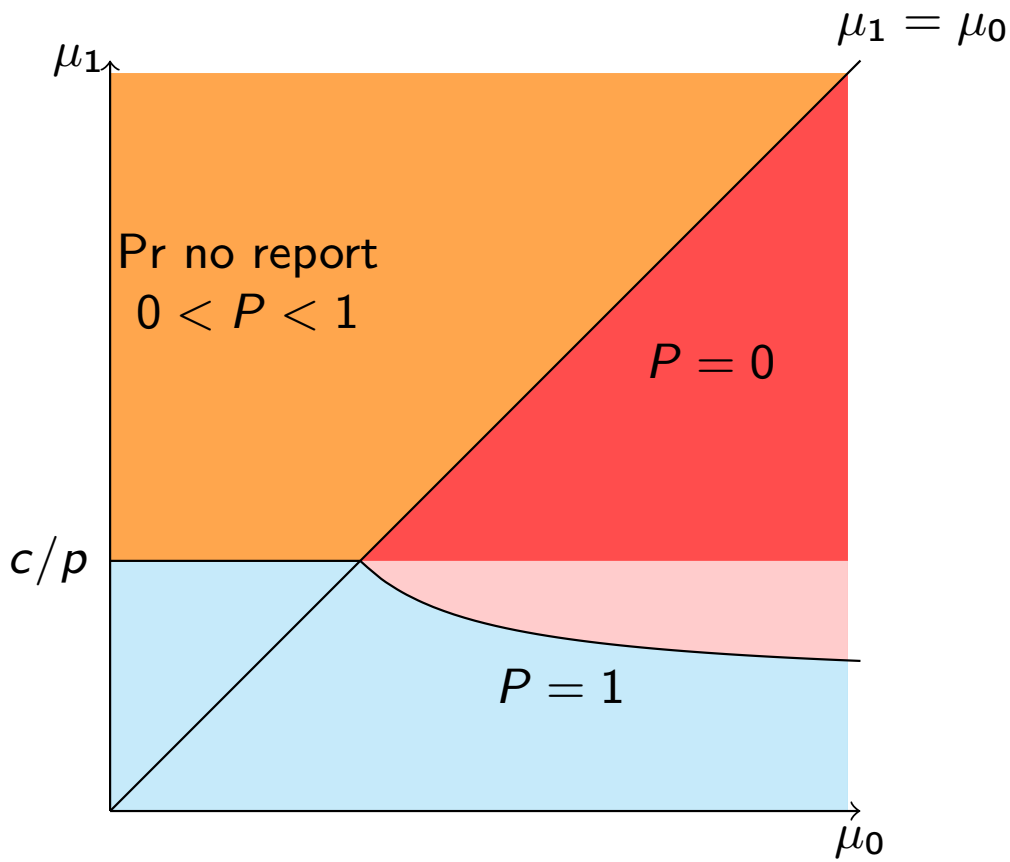

Optimal Protection

Theorem

If $\mu_0 \leq c/p$, then $\mu_1^*(\mu_0, c, p) = c/p$.

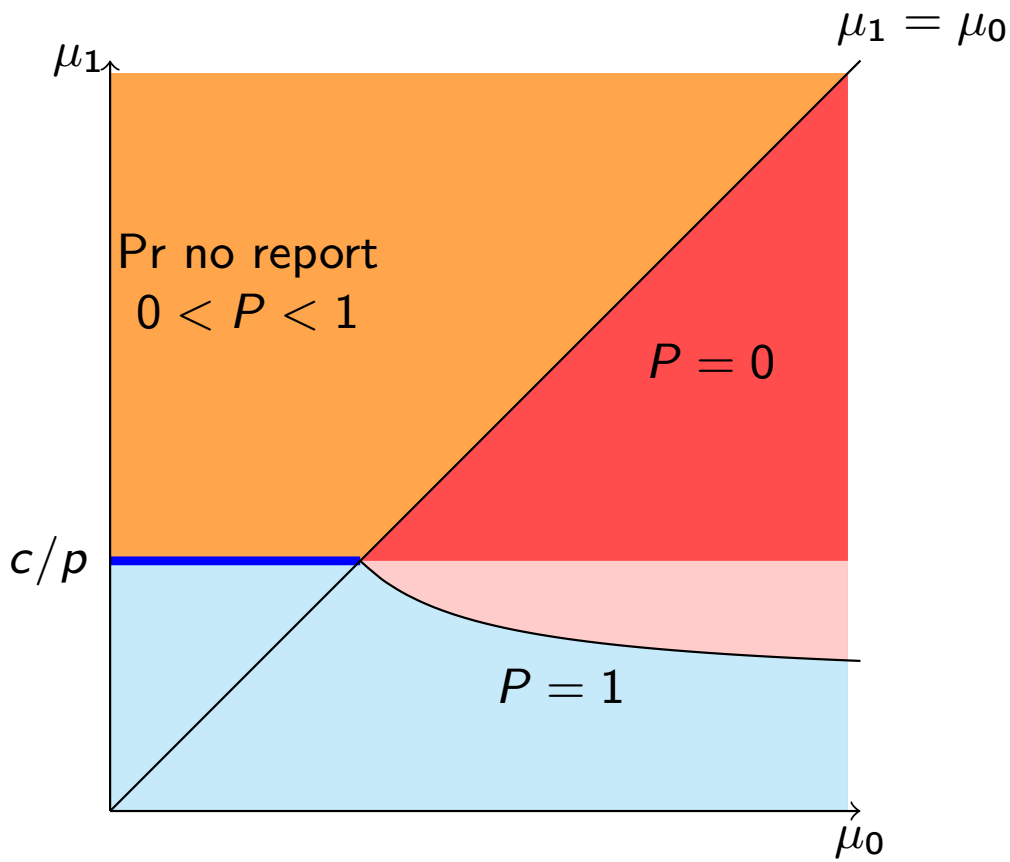

Optimal Protection

Theorem

If $\mu_0 \leq c/p$, then $\mu_1^*(\mu_0, c, p) = c/p$. If $\mu_0 > c/p$, then $\mu_1^*(\mu_0, c, p) > \mu_0$ uniquely solves

$$\underbrace{\frac{\mu_1 T(\mu_1, \mu_0, c, p)}{\mu_1 - \mu_0}}_{\text{speed effect}} = \underbrace{T(\mu_1, \mu_0, c, p)}_{\text{revelation effect}} + \underbrace{\frac{1}{(\mu_1 - \mu_0)(1 - c/\mu_1)}}_{\text{cost effect}}.$$

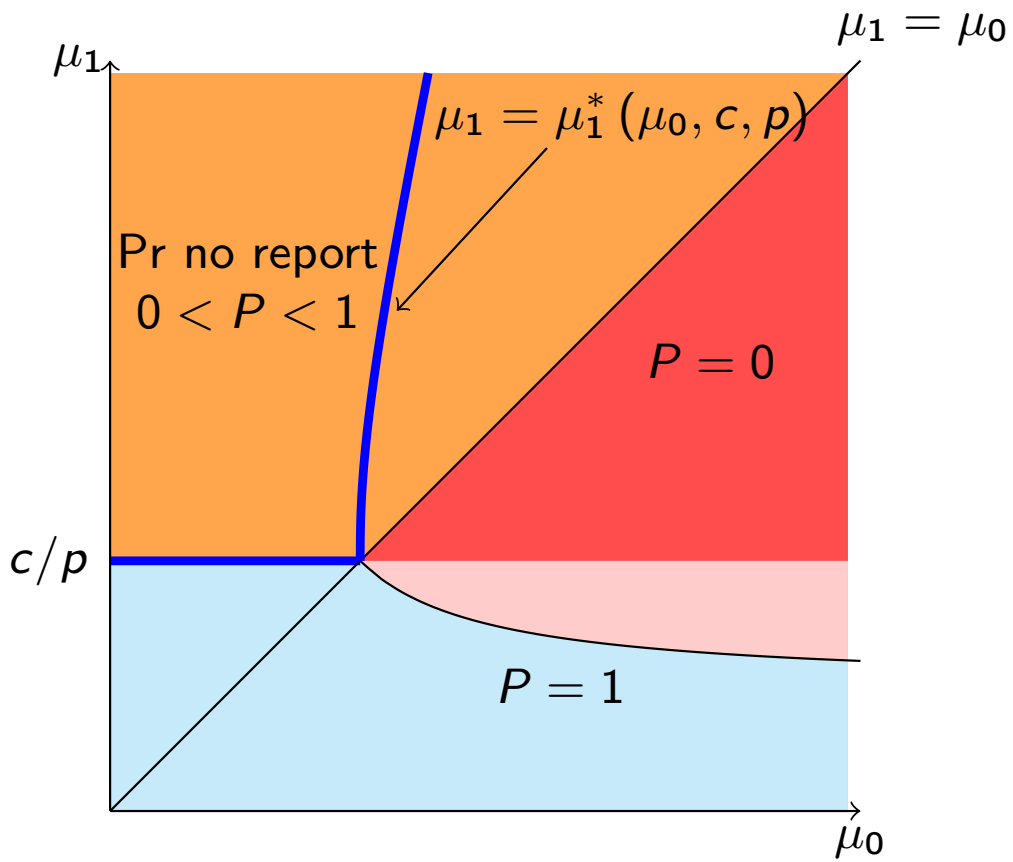

Result 2

The hider benefits from lower μ_0, that is, from stronger default protection.

Result 2

The hider benefits from lower μ_0, that is, from stronger default protection.

- intuition: $\downarrow \mu_0 \Rightarrow \uparrow (\mu_1 - \mu_0) \Rightarrow \uparrow$ speed of learning $\Rightarrow$ $\downarrow$ duration of learning $\Rightarrow \uparrow$ Pr(no report)

Result 2

The hider benefits from lower μ_0, that is, from stronger default protection.

- intuition: $\downarrow \mu_0 \Rightarrow \uparrow (\mu_1 - \mu_0) \Rightarrow \uparrow$ speed of learning $\Rightarrow$ $\downarrow$ duration of learning $\Rightarrow \uparrow$ Pr(no report)

Privacy Debate

- default protection — government responsibility
- selective protection — individual responsibility

Goal Protection of sensitive (compromising) information

Problem Default protection has high indirect cost since it limits access to big data

Question Could providing tools for selective protection be a solution?

Result 2

The hider benefits from lower μ_0, that is, from stronger default protection.

- intuition: $\downarrow \mu_0 \Rightarrow \uparrow (\mu_1 - \mu_0) \Rightarrow \uparrow$ speed of learning $\Rightarrow$ $\downarrow$ duration of learning $\Rightarrow \uparrow$ Pr(no report)

Privacy Debate

- default protection — government responsibility
- selective protection — individual responsibility

Goal Protection of sensitive (compromising) information

Problem Default protection has high indirect cost since it limits access to big data

Question Could providing tools for selective protection be a solution?

Answer No! Tools that facilitate selective protection might not be used in practice in the absence of strong default protection

Outline

Multiple Seekers

- Hider cannot control μ_1 but can choose the number of seekers $n \geq 1$.

Multiple Seekers

- Hider cannot control μ_1 but can choose the number of seekers $n \geq 1$.
- Private learning: the seekers do not communicate with each other

Multiple Seekers

- Hider cannot control μ_1 but can choose the number of seekers $n \geq 1$.
- Private learning: the seekers do not communicate with each other
- Competition: only the first seeker who reports a compromising story derives a positive payoff (reports are public)

Multiple Seekers

- Hider cannot control μ_1 but can choose the number of seekers $n \geq 1$.
- Private learning: the seekers do not communicate with each other
- Competition: only the first seeker who reports a compromising story derives a positive payoff (reports are public)

Revelation effect $\uparrow n \Rightarrow \uparrow$ Pr story is revealed when the seekers are searching for it $\Rightarrow \downarrow$ Pr(no report)

Speed effect $\uparrow n \Rightarrow p(t) \downarrow$ faster $\Rightarrow \uparrow$ Pr(no report)

Multiple Seekers

- Hider cannot control μ_1 but can choose the number of seekers $n \geq 1$.
- Private learning: the seekers do not communicate with each other
- Competition: only the first seeker who reports a compromising story derives a positive payoff (reports are public)

Revelation effect $\uparrow n \Rightarrow \uparrow$ Pr story is revealed when the seekers are searching for it $\Rightarrow \downarrow$ Pr(no report)

Speed effect $\uparrow n \Rightarrow p(t) \downarrow$ faster $\Rightarrow \uparrow$ Pr(no report)

No cost effect n does not change the stopping belief threshold $p(T)$

Multiple Seekers

- Hider cannot control μ_1 but can choose the number of seekers $n \geq 1$.
- Private learning: the seekers do not communicate with each other
- Competition: only the first seeker who reports a compromising story derives a positive payoff (reports are public)

Revelation effect $\uparrow n \Rightarrow \uparrow$ Pr story is revealed when the seekers are searching for it $\Rightarrow \downarrow$ Pr(no report)

Speed effect $\uparrow n \Rightarrow p(t) \downarrow$ faster $\Rightarrow \uparrow$ Pr(no report)

No cost effect n does not change the stopping belief threshold $p(T)$

Theorem

Under assumption $\mu_1 \geq c/p$ and no discounting, the optimal privacy protection is $n = +\infty$, that is, open access is optimal.

Multiple Seekers

- Hider cannot control μ_1 but can choose the number of seekers $n \geq 1$.
- Private learning: the seekers do not communicate with each other
- Competition: only the first seeker who reports a compromising story derives a positive payoff (reports are public)

Revelation effect $\uparrow n \Rightarrow \uparrow$ Pr story is revealed when the seekers are searching for it $\Rightarrow \downarrow$ Pr(no report)

Speed effect $\uparrow n \Rightarrow p(t) \downarrow$ faster $\Rightarrow \uparrow$ Pr(no report)

No cost effect n does not change the stopping belief threshold $p(T)$

Theorem

Under assumption $\mu_1 \geq c/p$ and no discounting, the optimal privacy protection is $n = +\infty$, that is, open access is optimal.

Theorem (Result 3)

Under assumption $\mu_1 > c/p$ and discount rate $\rho > 0$, there exists $\rho^* > 0$ such that open access ($n = +\infty$) is optimal if $\rho < \rho^*$ and the strongest feasible protection ($n = 1$) is optimal if $\rho > \rho^*$.

Outline

Literature

- Signaling games: *Hagenbach and Koessler (2017), Daughety and Reinganum (2010), Gradwohl and Smorodinsky (2017)*
  - action is observable; hider chooses suboptimal action (such as low protection) to avoid signaling
  - in our model, hider's realized level of protection is unobservable

Literature

- Signaling games: *Hagenbach and Koessler (2017), Daughety and Reinganum (2010), Gradwohl and Smorodinsky (2017)*
  - action is observable; hider chooses suboptimal action (such as low protection) to avoid signaling
  - in our model, hider's realized level of protection is unobservable
- Privacy as a tool:
  - consumer privacy where anonymity helps to avoid price discrimination: *Fudenberg and Villas-Boas (2006), Calzolari and Pavan (2006), Conitzer at al (2012)*
  - reputation models where hider cares about public belief about her either per se (*Daughety and Reinganum (2010)*) or in context of repeated games (*Mailath and Samuelson (2006)*)
  - in our model, privacy has an intrinsic value
    - privacy in *Gradwohl (2018), Dziuda and Gradwohl (2015), Gradwohl and Smorodinsky (2017)* also has intrinsic value but they study different environments

Literature

- Information design literature: *Kamenica and Gentzkow (2011), Bergemann and Morris (2019)*
  - we borrow: commitment by an uninformed discloser

Literature

- Information design literature: *Kamenica and Gentzkow (2011), Bergemann and Morris (2019)*
  - we borrow: commitment by an uninformed discloser
  - we differ: restrictive set of feasible disclosure rules
    - as in advertising literature (*Johnson and Myatt (2006), Anderson and Renault (2006)*), literature on auction design (*Milgrom and Weber (1982)*) and literature on second-degree price discimination (*Ottavianni and Prat (2001)*)

Literature

- Information design literature: *Kamenica and Gentzkow (2011), Bergemann and Morris (2019)*
  - we borrow: commitment by an uninformed discloser
  - we differ: restrictive set of feasible disclosure rules
    - as in advertising literature (*Johnson and Myatt (2006), Anderson and Renault (2006)*), literature on auction design (*Milgrom and Weber (1982)*) and literature on second-degree price discimination (*Ottavianni and Prat (2001)*)
- Strategic experimentation with Poisson bandits: *Keller et al (2005)*
  - trade-off between effectiveness and duration of learning: *Bobtcheff and Levy (2017), Halac et al (2017), Cetemen and Margaria (2021)*

Literature

- Information design literature: *Kamenica and Gentzkow (2011), Bergemann and Morris (2019)*
  - we borrow: commitment by an uninformed discloser
  - we differ: restrictive set of feasible disclosure rules
    - as in advertising literature (*Johnson and Myatt (2006), Anderson and Renault (2006)*), literature on auction design (*Milgrom and Weber (1982)*) and literature on second-degree price discimination (*Ottavianni and Prat (2001)*)
- Strategic experimentation with Poisson bandits: *Keller et al (2005)*
  - trade-off between effectiveness and duration of learning: *Bobtcheff and Levy (2017), Halac et al (2017), Cetemen and Margaria (2021)*
  - private learning with many agents: *Akcigit and Liu (2016)* where players do not share bad news

Outline

Conclusion

- Given the default protection for non-compromising information, the hider commits to the selective protection for compromising information.

Conclusion

- Given the default protection for non-compromising information, the hider commits to the selective protection for compromising information.
- If the default protection is weak, the optimal selective protection is weak
  - hider trades the effectiveness of the seeker's learning and its duration

Conclusion

- Given the default protection for non-compromising information, the hider commits to the selective protection for compromising information.
- If the default protection is weak, the optimal selective protection is weak
  - hider trades the effectiveness of the seeker's learning and its duration
- The hider benefits from stronger default protection $\Rightarrow$ we need strict privacy laws