08375
Economics of Uncertainty and Information

Topic 3: Information

Plan:

- Topic 1: Introduction and Recap
- Topic 2: Uncertainty
- Topic 3: Information
- Topic 4: Current Research Areas

Information

Moral Hazard

Plan:

- Topic 3: Information
  - 3.1 Adverse Selection
    - Lemons
    - Other applications
    - Experimental evidence
  - 3.2 Signalling and Screening
    - Signalling in the job market
    - Screening in the job market
    - Other applications
    - Experimental evidence
  - 3.3 Moral Hazard
    - Moral hazard in the insurance market
    - Other applications
    - Experimental evidence

Topic Outline

- By the end of this sub-topic, you should be able to analyse the following:
  - Describe the problem of hidden action and how this can lead to moral hazard
  - Analyse the shareholder-manager problem intuitively and algebraically
  - Describe and explain the role of risk preferences in the principal-agent problem
  - Describe, interpret and critically assess the market experiment by DeJong, Forsythe, Lundholm and Uecker (1985)

What's it all about?

- So far we've considered cases where agents had an incentive to **hide some information**
  - They took advantage of **pre-contractual opportunism**
  - The result was **inefficiency**, with over- or under-employment, foregone gains from trade and low quality trade (if trade occurred at all).
- Now we consider cases where agents have an incentive to **hide their actions**
  - They take advantage of **post-contractual opportunism**
  - Their actions cannot be observed.
  - After they agree to act in a certain way, they can cheat and act differently instead.
- We are interested in cases where one party's **incentives are misaligned** with the incentives of the other party. There is **moral hazard**.

What's it all about?

- Cheating
  - Tests
  - Tax returns
  - Card games
  - Drugs in sports
  - Romantic relationships
  - Cartels
  - Shirking workers
  - Broken promises

"Any situation in which one person makes the decision about how much risk to take, while someone else bears the cost if things go badly."

Krugman, Paul (2009). The Return of Depression Economics and the Crisis of 2008. W.W. Norton Company Limited. ISBN 978-0-393-07101-6.

What's it all about?

- If you buy health insurance, are you more or less likely to smoke?

- If a bank is seen as too big to fail, what will be its incentive to make safe investment decisions?

What's it all about?

- To put some structure on the problem, first we consider **principal-agent** problems.
- One party is the **principal**.
- Another party is the **agent**
- The agent acts on behalf of the principle, but may have different aims and objectives.
- If the principal cannot observe the agent's actions, the agent may not act in the best interests of the principal, even if they are contracted to do so.

What's it all about?

- The principal cannot observe or verify the action of the agent
- The principal can observe the eventual outcome of the agent's action
- For example, the agent could put in high effort or low effort on a project. If he puts high effort in, the project is more likely to succeed. But it is *possible* that he could succeed even if he puts in low effort, or he could fail even if he puts in high effort.
- Observing the outcomes gives the principal a clue about the action, but it is not a perfect signal.
- Putting in low effort has consequences for the principal so the agent does not bear full consequences of this low effort.

What's it all about?

- We will explore how this set-up plays out in the context of managerial and shareholder decision-making.
- We will then explore some other applications of moral hazard
- Finally we will look briefly at some evidence from experiments and the real world more generally.

Managerial decision making

- Principal: shareholder. Wishes to maximise profits

- Agent: manager. May have differing aims:
  - Preference for perks and expenses
    - Claims to need a company car when it is not necessary.
  - Preference for leisure
    - Claims to work hard but really shirks
  - Preference for lower risk actions
    - If a project fails the manager may lose their position. This may make them more risk averse than the shareholders, especially if shareholder portfolios are diversified
  - Time preference
    - Managers may prefer short run successes over larger successes in the future

Managerial decision making

- Adding some technical detail…
- Shareholder (principal)
  - Maximises expected net profits: $\max E(\pi - w)$
    - π is gross profit from the business activity, and can be high (π_H) or low (π_L)
    - w is remuneration paid to the agent
  - Assumed to be risk neutral due to a diversified portfolio
- Manager (agent)
  - Maximises expected utility from wages
  - Assumed to be risk averse
  - Can choose to put in effort (e_H) or shirk (e_L).

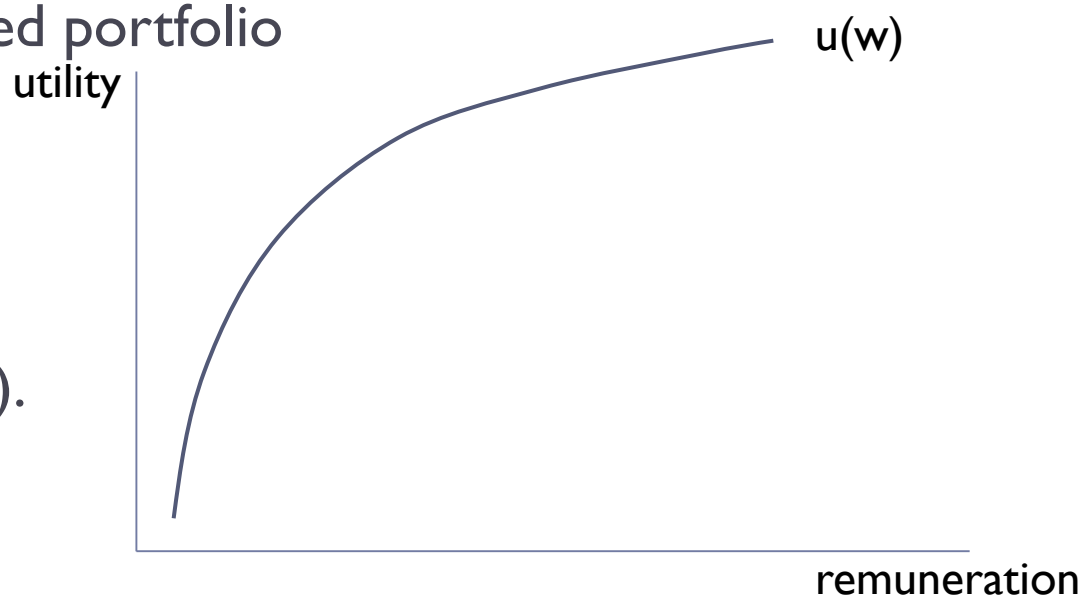

Managerial decision making

- The success of the project depends on the level of effort put in by the manager.
- However, the success is not deterministically related to effort. There is also a role for chance.
  - Even with high effort, the project may fail
    - So failure cannot be interpreted as definitely revealing low effort
  - Even with low effort, the project may succeed
    - So success cannot be interpreted as definitely revealing high effort

- Probabilities x and y give the chance of success (π_H) instead of failure (π_L), conditional on high and low effort, respectively.

Managerial decision making

- On a diagram

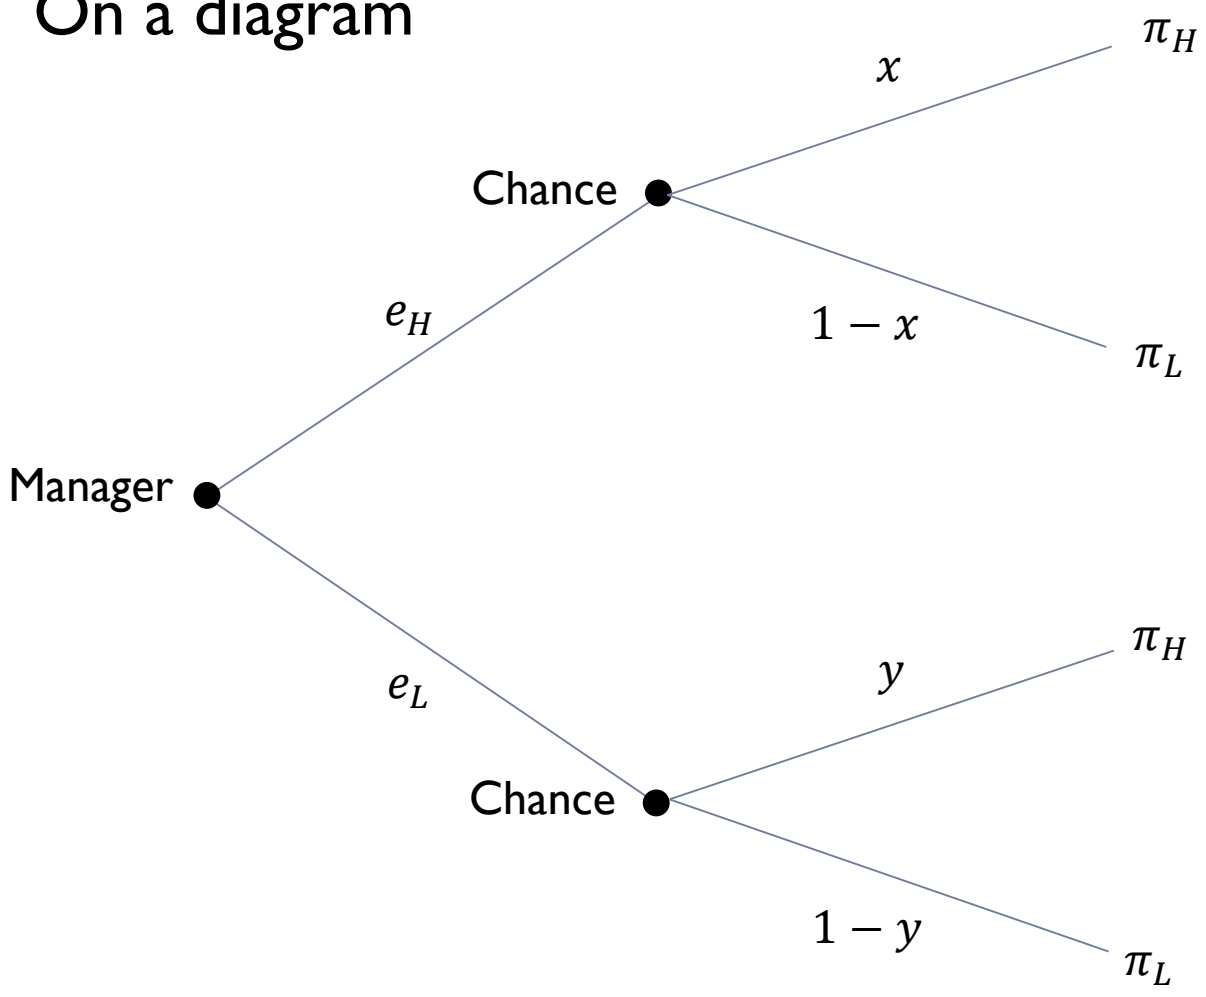

Managerial decision making

- Let the manager have a willingness to accept compensation for hard work that is given by $\emptyset$.
  - If the wage difference between hard work versus shirking is at least $\emptyset$ then the manager will work hard. If not they will shirk.

- Specifically, they decide based on their utility:
  - If work hard (e_H): utility is given by $U(w - \emptyset)$
  - If shirk (e_L): utility is given by $U(w)$

- Let there be an outside option of working for w_0 so $U_0 = u(w_0)$

Managerial decision making

- The shareholder decides the remuneration of the manager
- Two basic constraints:
  - $U \geq u(w_0)$
  - $(\pi - w) \geq 0$
- If the constraints are not satisfied then either the managers or the shareholders withdraw from the agreement.

- But there is more to say about the structure of the remuneration.
  - First with full information
  - Then with asymmetric information and hence Moral Hazard

Managerial decision making: full information benchmark

- First, assume **full and verifiable** information about manager's effort level.
- Shareholder selects optimal level of effort from the manager to maximise net profits
- Since manage is risk averse, the best the shareholder can do is offer a fixed wage independent of the success of the project.
  - Do not need to pay a risk premium because there is no risk being faced.
  - Wage will be just high enough that the manager won't leave.
- This gives us **optimal risk sharing** because the **risk neutral** shareholder faces all of the risk, and the **risk averse** manager faces no risk at all.

- **But what will the flat wage actually be?**

Managerial decision making: full information benchmark

- Two relevant cases:

Case A:

- Shareholder demands low effort and pays flat wage exactly equal to w_0 (outside wage option).
- $U_0 = U(w_0)$

Case B:

- Shareholder demands high effort and pays flat wage equal to $w_0 + \emptyset$ (outside wage option plus min WTA for high effort).
- $U_1 = U([w_0 + \emptyset] - \emptyset)$
- This means $U_1 = U_0$

Managerial decision making: full information benchmark

- Whether Case A or Case B holds depends on the relative probabilities of success, and on the size of the high effort compensation

- Specifically, Case B will happen if:
  - $x\pi_H + (1-x)\pi_L - \emptyset > y\pi_H + (1-y)\pi_L$
  - $(x-y)\pi_H + (\pi_L - \pi_L) - (x+y)\pi_L - \emptyset > 0$
  - $(\boldsymbol{x}-\boldsymbol{y})(\boldsymbol{\pi_H} - \boldsymbol{\pi_L}) > \emptyset$

- We have outlined the full information case.
  - High effort and wage = $w_0 + \emptyset$ if $(x-y)(\pi_H - \pi_L) > \emptyset$
  - Low effort and wage = w_0 if $(x-y)(\pi_H - \pi_L) \leq \emptyset$

Managerial decision making: hidden action

- What if the shareholder cannot observe the effort made by the manager?

- In Case A: there is no problem!
  - The shareholder wants low effort and pays w_0.
  - The manager can do no better than to put in low effort and receive w_0
  - No moral hazard arises.

- However, in Case B it is more complicated.
  - The shareholder wants high effort
  - But the manager can claim to put in high effort, but actually put in low effort instead.
  - If they get away with it, they get the wage = $w_0 + \emptyset$ and utility $u(w_0 + \emptyset) > u(w_0)$
  - If the project fails, they can claim it was due to chance, not low effort.

Managerial decision making: hidden action

- This is Moral Hazard.
- So what can the shareholder do?

- Can they link pay to EFFORT?
  - No! It is not observable

- Can they link pay to PROFIT?
  - Yes! Can link the pay to the gross profit (π_H or π_L)
  - Bonuses, profit-related pay, profit-sharing
- We seek 2 wages, w_H to paid when $\pi = \pi_H$ and w_L to be paid when $\pi = \pi_L$
- First, set the **constraints**

Managerial decision making: hidden action

<u>Incentive compatibility constraint</u>

- The shareholder must set w_L and w_H such that:
  - $EU_{work} \geq EU_{shirk}$ i.e. $EU_{e_H} \geq EU_{e_L}$

- $EU_{e_H} = xU(w_H - \emptyset) + (1 - x)U(w_L - \emptyset)$
- $EU_{e_L} = yU(w_H) + (1 - y)U(w_L)$

- It must be in the manager's own best interests to work hard.
- The manager's **incentives** must be **compatible** with the desires of the shareholder

Managerial decision making: hidden action

<u>Participation constraint</u>

- The shareholder must set w_L and w_H such that:
  - $EU_{work} \geq U_0$

- It must be in the manager's own best interests to work instead of taking the outside option.
- The manager must be incentivised to **participate** in the employment.

Managerial decision making: hidden action

<u>Solution for net profit maximisation</u>

- The weak inequalities can be replaced by strict equalities because setting wages higher is *wasteful* from the perspective of the shareholder.
  - E.g. can set wages conditional on success to be astronomically high. The employee will still participate. But it won't maximise shareholder profits!
- Therefore:
  - Incentive compatibility constraint gives $EU_{work} = EU_{shirk}$
  - Participation constraint gives $EU_{work} = U_0$
- Together:
  - $\boldsymbol{EU_{work} = EU_{shirk} = U_0}$
- **Intuition**: shareholder sets wages **as low as possible** on average to give max net profit, constrained by need for manager to **participate** and **work hard**

Managerial decision making: hidden action

<u>Features of the solution</u>

- Risk-sharing is no longer optimal
  - Risk averse manager now faces some uncertainty over their wages
  - This is necessary to satisfy incentive compatibility
  - But sub-optimal because the shareholder is risk neutral and ideally would absorb the risk
- Manager is equally well off compared to full information
  - Despite facing risk, the wage is higher in expectation than w_0 to compensate
  - It results from the participation constraint which binds in both full and imperfect information cases
- Shareholder is worse off compared to full information
  - The manager receives some compensation for facing uncertainty (risk premium) at the expense of the shareholder.

Managerial decision making: hidden action

<u>Detailed comparison between full and imperfect information</u>

- Expected wage payment under imperfect information with hard work is:
  - $xw_H + (1 - x)w_L$
- Wage under full information
  - $w_0 + \emptyset$
- Since the manager requires compensation for facing uncertainty:
  - $xw_H + (1 - x)w_L > w_0 + \emptyset$
  - The risk premium is $xw_H + (1 - x)w_L - w_0 - \emptyset$
- Since wages are higher under imperfect information, this means profits are lower and shareholders are worse off. Since managers are no better off, we can see that full information Pareto dominates imperfect information.
- The inefficiency arises due to sub-optimal risk sharing

Managerial decision making: hidden action

Alternative solutions

- Three options:
  - Set incentives that lead to e_H
    - Let net profits for shareholder be given by Π_{work}
  - Set incentives that lead to e_L
    - Let net profits for shareholder be given by Π_{shirk}
  - Shut down the firm altogether
    - Let net profits for shareholder be $\Pi_{shutdown} = 0$
- Which option will the shareholder choose?
  - It depends on whether Π_{work}, Π_{shirk} or $\Pi_{shutdown}$ is largest.
  - This depends on whether the profits from high effort are sufficiently high to offset the costs of inducing high effort; and whether the profits from low effort are negative or not.

Managerial decision making: hidden action

<u>Alternative solutions</u>

- The manager is ALWAYS equally well off under each option
  - Earn U_0 in either case, and if shut down then earn U_0 in outside employment

- The shareholder is ALWAYS worse off compared to full information
  - If induce high effort, incur cost of risk premium
  - If tolerate low effort, incur cost of foregone profits
  - If shut down, incur losses from foregone production

Managerial decision making: hidden action

Risk attitude

- We required a Risk Averse manager and a Risk Neutral principal in order to get the results we just discussed
- If the manager were also risk neutral then it would not be sub-optimal to share the risk.
- In this case, the principal (shareholder) could make the agent (manager) bear all of the risk.
- There would be no need for a risk premium to be paid so this would not be any more costly for the shareholder compared to paying a flat wage
- Therefore, the principal's profits would not suffer and there would essentially be no moral hazard problem.

Managerial decision making: summary

- With full information, optimal risk sharing means the manager should get a flat wage and the shareholder bear all the risk. The shareholder may choose the wage to induce high effort or low effort, depending on the cost of inducing high effort. They may even leave the market altogether.
- With imperfect information, the agent can pretend to do high effort and actually do low effort. This is possible because their effort is an imperfect determinant of eventual profits.
- With imperfect information, the shareholder may force some of the risk to be borne by the manager. This provides incentive compatibility. However, it is costly as the manager must be compensated (paid a risk premium).
- This cost may be too great, and so shareholders may settle for the low effort outcome.

Experimental evidence

- DeJong, Forsythe, Lundholm and Uecker (1985)

- General idea: principal contracts an agent, agent can then cheat or fulfil the promise. Final outcome depends on the agent's chosen action and on chance.

- Also test possible SOLUTIONS to moral hazard:
  - Multiple periods (possible reputation)
  - (Optional) public investigations of agent's action when bad result happens
  - (Optional) private investigations of agent's action when bad result happens
  - Agent bearing liability for cheating (if found out)
  - Agent privately bearing liability for cheating (if found out)

Experimental evidence

DFLU experimental set-up

- University of Iowa students
- 2.5 – 3 hour experimental s3essions
- Cash incentives from $10 - $25

Markets

- 3 agents endowed with $0.50
- 3 principals endowed with $1.50. They face a sure loss of $0.80 unless they employ an agent to provide a service
- 3 "quality levels" of service
  - Level 1 (Low): probability principal loses $0.80 = 0.8
  - Level 3 (Low): probability principal loses $0.80 = 0.5
  - Level 5 (Low): probability principal loses $0.80 = 0.1

Experimental evidence

<u>Prices</u>

- Prices are determined in a "sealed bid auction"
  - Step 1. Agents offer to provide a specified service quality for a specified price (or higher)
  - Step 2. Principals say which agent they'd like to enter into a contract with
  - Step 3. The match is announced (but not the price or the quality agreed upon)
  - Step 4. The agent decides what quality they will really provide. They pay the associated cost.
  - Step 5. The relevant probability (depending on the ACTUAL action chosen) is resolved.
  - Step 6. The agent and principal find out whether the loss occurs or not.
  - {here, some additional steps on investigation and liability depending on the treatment}
- Then game starts again from Step 1.
- Earnings accumulate until the final round, when they are cashed in.

Experimental evidence

Details of the quality services

Principal			Agent	Total	
Level of service q	Probability of loss $P(q)$	Expected loss $P(q)l$ (\$)	Cost of service $c(q)$ (\$)	Expected cost (\$)	Expected surplus (\$)
No purchase	1	0.8	0	0.8	1
1	0.8	0.64	0.08	0.72	1.08
3	0.5	0.40	0.2	0.60	1.20
5	0.1	0.08	0.65	0.73	1.07

Level 3 quality gives greatest total expected surplus

$P(q)$ given on the previous slide

$P(q)$ multiplied by the loss of 0.8

Cost to agent of providing service

Sum of $P(q)l$ and $c(q)$

Principal's expected return plus agent's expected return:
$(\$1.30 - price - P(q)l)$
$+ (\$0.50 + price - c(q)$
$= \$1.80 - total\ expected\ cost$

Experimental evidence

Predictions

- Benchmark (full information)
  - Principals can observe and verify the quality of service
  - Under risk neutrality (strong assumption!):
    - Maximise expected net profit
    - $\max(E(\pi) = \$1.30 - price\ - P(q) \cdot l$
  - With perfect competition (strong assumption again!):
    - Zero profits in equilibrium
    - $price = c(q)$
  - Combining these statements:
    - $\max(E(\pi) = \$1.30 - c(q) - P(q) \cdot l$
    - This is maximised where the quality level is 3 (i.e. medium) – see table on previous slide.
- In the benchmark, trade will all be in quality level 3 and price = $0.20

Experimental evidence

<u>Predictions</u>

- Moral hazard in one period case
  - Principals cannot observe and verify the quality of service
  - Agents will therefore always provide the lowest quality service to minimise their costs
    - Any higher quality of service is a dominated strategy in a one-shot game.
  - With perfect competition (strong assumption again!):
    - Zero profits in equilibrium
    - $price = c(q)$
    - $price = c(1) = 0.08$
- In the one-shot moral hazard case trade will all be in quality level 1 and price = $0.08

Experimental evidence

Predictions

- Moral hazard in one period case plus identification (checking)
  - There is no benefit of identifying cheaters because there is no future round for retaliation
  - Standard theory predicts NO difference compared to moral hazard in the one-shot case without identification.
- Moral hazard in one period case plus identification (checking) AND liability
  - If those agents caught cheating have to compensate the principal for a resulting loss, then the equilibrium changes.
  - DFLU prove a mixed strategies nash equilibrium
    - $price = \$0.20; \alpha = 0.4; \rho = 0.19$
      - ρ is probability that the principal (pays to) investigate
      - α is probability that the agent cheats (q=1)

Experimental evidence

Predictions

- Multiple rounds (reputation)
  - No sharp theoretical predictions arise due to multiple possible equilibria
  - Four hypotheses (from DFLU):
    - With multiple periods, there is non-zero frequency of investigation even in the absence of the liability rule because principals can form expectations about whether agents tend to cheat based on outcomes observed over time
    - With multiple periods, there is less incidence of cheating even in the absence of the liability rule because agents have an incentive to preserve their reputation and be paid for contracts of level 3.
    - Private disclosure leads to more investigations because the principals can't free ride on others' investigations.
    - Public disclosure leads to less cheating than private disclosure.

Experimental evidence

Numbers in the table = average over the whole experiment
Numbers in brackets = average over only last 5 rounds

Results

Market set-up	Market	Frequency of $q = 1$ delivery (α)	Frequency of investigation (ρ)	price	Efficiency
'No investigation, no liability'	1	0.53 (0.41)	-	0.33 (0.32)	0.54 (0.53)
	2	0.68 (0.90)	-	0.26 (0.20)	0.51 (0.43)
	Average (1,2)	**0.61 (0.68)**		**0.29 (0.26)**	**0.53 (0.48)**
'Investigation, no liability'	3	0.47 (0.88)	0.11 (0)	0.28 (0.23)	0.55 (0.55)
	4	0.20 (0.37)	0.06 (0.07)	0.39 (0.31)	0.64 (0.67)
	Average (3,4)	**0.39 (0.61)**	**0.09 (0.04)**	**0.31 (0.27)**	**0.59 (0.55)**
'Investigation and liability' (public)	5	0.04 (0.13)	0.49 (0.17)	0.46 (0.35)	0.38 (0.58)
	6	0.03 (0.05)	0.19 (0)	0.32 (0.23)	0.73 (0.97)
	Average (5,6)	**0.04 (0.09)**	**0.32 (0.09)**	**0.39 (0.29)**	**0.55 (0.76)**
'Investigation and liability' (private)	**7**	**0.20 (0.15)**	**0.61 (0.75)**	**0.33 (0.33)**	**0.26 (0.16)**

Experimental evidence

<u>Results</u>

- Some evidence of $q = 5$ which is not predicted by any theory
- 'baseline' case with no liability or investigation shows around 60-70% delivering low quality as predicted, 30-40% delivering higher quality despite the predictions. Price was higher than predicted ($0.30 instead of $0.08)
- 'investigation' gives a big drop in cheating. Around 1 in 10 losses are investigated. Prices were not much higher, despite the better quality delivered.
- 'liability' is powerful with rates of cheating dropping to close to 0 and prices substantially higher than in the other cases
- Private investigation and liability was much less effective than public

Experimental evidence

Results

- Efficiency is calculated as the proportion of total possible surplus that was actually captured

- $$efficiency = \frac{average\ realised\ trading\ surplus\ per\ period}{maximum\ potential\ expected\ surplus\ per\ period}$$

- Very little difference between markets 1- 6 but very poor surplus in market 7 with private investigations and liability.

Experimental evidence

Conclusions of DFLU

- Cheating happens
- Reputation-building appears to happen (investigation even without liability)
- Delivery above minimum quality happens even without investigation or liability, suggesting possible reputation effects
- Public disclosure of cheating is somewhat powerful in reducing cheating but this is strengthened by liability rules
- Private disclosure is inefficient
- Not very efficient overall!

To sum up:

- Hidden action can lead to moral hazard
- Agents' **incentives are misaligned** with the incentives of the principal.
- The principal bears some of the cost of the downside risks taken by the agent
- We studied a manager-shareholder framework and showed how shareholders can set wages to realign their agent's incentives with their own
- We showed how the final position is equally good for the agent but worse for the principal
- We discussed an experimental test in detail (DFLU)

References and readings:

- DeJong, Forsythe, Lundholm and Uecker (1985) "Ripoffs, lemons and reputation formation in agency relationships: a laboratory market study" *Journal of finance*, 50, pp 809-20