EC334 Topics in Financial Economics
Third Seminar

Junxi Liu

February 2024

Corporate Policy

- CWS13: The Rolf of the CFO, Performance Measurement, and Incentive Design
  - Read the book
- CWS15: Capital Structure and the Cost of Capital: Theory and Evidence
  - Equity versus Debt to fund a firm
  - Value of a Firm
  - Modigliani-Miller propositions
- CWS16: Dividend Policy: Theory and Empirical Evidence
  - How to pay dividend without and with taxes
  - Repurchases

Financial theory and corporate policy
by Copeland, Thomas E; Weston, J. Fred; Shastri, Kuldeep

Imagine You Are Running a Firm…

- Your firm needs money.
  - (Unless you are from a billionaire family)
- You can borrow money.
  - Debt: requires repayment; often with interests; usually tax-free
- You can also sell some shares of your firm
  - Equity: an ownership in the firm; no need to repay; if your firm skyrockets then you are losing a LOT of money…
- Gearing (or Leverage, in the rest of the world)
  - In simple words: borrowing money to invest
  - For corporate policy: the relationship, or ratio, of a company's debt-to-equity
  - For example, a gearing ratio of 70% shows that a company’s debt levels are 70% of its equity.

Imagine You Are Running a Firm…

- You have to keep a balance of debt and equity
- If you know your business is going to be a great success but the investors in the market don’t think so
  - you will only be able to sell your shares (equity) at a very low price — which you believe can be a LOT of money in the future!
- So you want to borrow money instead — you don’t have to “sell yourself”.
  - However, borrowing means that you will have to repay the money within a certain time frame, and also pay quite a lot interests — you are under pressure!

Let’s See an Successful Story of Equity Investment

- BYD was the world's largest plug-in hybrids and pure electric vehicles maker in 2022, with a total of 1.86 million cars sold.

ADL BYD to build 130-strong ZE fleet for National Express in Coventry

By **routeone Team** - February 17, 2022

arriva
a DB company

Careers

Arriva – Network Manager

routeone Team - February 22, 2023

Arriva are looking to recruit a Network Manager to b
a part of its North Midlands operation. Network Manager North Midlands Arriva's vision is to shape a...

East Yorkshire Buses announces plans for new Scarborough depot

February 22, 2023

**First Bus leaves Southampto
as Go-Ahead steps in**

Let's See an Successful Story of Equity Investment

- In late 2008, Berkshire Hathaway ponied up the aforementioned \$232 million for a roughly 10% stake in BYD.
- As Buffett recalled, Berkshire initially tried to buy 25% of the company, but Wang (CEO of BYD) refused to release more than 10% of BYD's stock.
- Currently, BYD's market capitalization is ~\$70 billion, which means a 10% stake would worth \$7 billion
- What if Wang used debt to finance, instead of selling the shares to Buffett?

When There Are Tax and Bankruptcy Costs…

- In many countries, interest is *deductible as a cost* of doing business while dividends are taxed as income – obviously favors debt financing
  - Tax shield of debt (affect Earnings Before Interest and Taxes EBIT)
  -

Interest Tax Shield Calculation				
($ in 000s)			Company A	Company B
Revenue			$50,000	$50,000
Less: Cost of Goods Sold (COGS)			(10,000)	(10,000)
Gross Profit			**$40,000**	**$40,000**
Less: Operating Expenses (OpEx)			(5,000)	(5,000)
EBIT			**$35,000**	**$35,000**
Less: Interest Expense			--	(4,000)
Pre-Tax Income (EBT)			**$35,000**	**$31,000**
Less: Taxes	Tax Rate	21.0%	(7,350)	(6,510)
Net Income			**$27,650**	**$24,490**
			Tax Shield	**$840**

- $PV = PV$(equity) + PV(tax shield) – PV(distress costs)
  - If you borrow way too much, investors will lose confidence and your firm might go into bankruptcy

Tax Shield

Assume for the moment that there are only two types of personal tax rates: the rate on income received from holding shares, τ_{ps}, and the rate on income from bonds, τ_{pB}. The expected after-tax stream of cash flows to shareholders of an all-equity firm would be $(EBIT)(1-\tau_c)(1-\tau_{ps})$. By discounting this perpetual stream at the cost of equity for an all-equity firm, we have the value of the unlevered firm:

$$V_U = \frac{E(EBIT)(1-\tau_c)(1-\tau_{ps})}{\rho}. \tag{21}$$

ρ → discount rate for an unlevered firm

Alternatively, if the firm has both bonds and shares outstanding, the earnings stream is partitioned into two parts. Cash flows to shareholders after corporate and personal taxes are

$$\text{payments to shareholders} = (EBIT - k_d D)(1-\tau_c)(1-\tau_{ps}),$$

and payments to bondholders, after personal taxes, are

$$\text{payments to bondholders} = k_d D(1-\tau_{pB}).$$

interest rate on bond; principal

$$\begin{array}{c}\text{total cash payments}\\ \text{to suppliers of capital}\end{array} = EBIT(1-\tau_c)(1-\tau_{ps}) - k_d D(1-\tau_c)(1-\tau_{ps}) + k_d D(1-\tau_{pB}). \tag{22}$$

Tax Shield

The first term on the right-hand side of (22) is the same as the stream of cash flows to owners of the unlevered firm, and its expected value can be discounted at the cost of equity for an all-equity firm. The second and third terms are risk free and can be discounted at the risk-free rate, k_b. The sum of the discounted streams of cash flow is the value of the levered firm:

$$\begin{aligned} V_L &= \frac{E(EBIT)(1-\tau_c)(1-\tau_{ps})}{\rho} + \frac{k_d D\left[(1-\tau_{pB}) - (1-\tau_c)(1-\tau_{ps})\right]}{k_b} \\ &= V_U + \left[1 - \frac{(1-\tau_c)(1-\tau_{ps})}{(1-\tau_{pB})}\right] B, \end{aligned} \tag{23}$$

where $B = k_d D(1-\tau_{pB})/k_b$, the market value of debt. Consequently, with the introduction of personal taxes, the gain from leverage is the second term in (23):

$$G = \left[1 - \frac{(1-\tau_c)(1-\tau_{ps})}{(1-\tau_{pB})}\right] B. \tag{24}$$

Discussion on the Firm Value with Personal Tax

- Essentially, the gain from having debt is $\left[1-\frac{(1-\tau_i^E)(1-\tau_c)}{1-\tau_i^B}\right]B$
  - When it is positive, firm has the incentive to hold debts
  - In addition, investor's demand for bonds changes with tax rate on interest income, and investors will be indifferent between bonds and equity if
$$r_D\left(1-\tau_i^D\right)=r_E\left(1-\tau_i^E\right)$$
  - Firms, who supply bonds, will be indifferent between supplying and not supplying if
$$r_D=\frac{r_E}{\left(1-\tau_C\right)}$$
  - Connecting the above two equations, you will get
$$\left(1-\tau_i^E\right)\left(1-\tau_c\right)=1-\tau_i^B$$
  - Which means tax shield is 0!
- Think about the real world…
  - In reality, it is often reasonable to argue that tax on stock is lower than that on bonds; to make bonds attractive firms have to pay more interests — tradeoff!
  - Gain to leverage is offset by the differential in bond tax and capital gain tax

One more word…

Rishi Sunak paid effective tax rate of 23% on £2.2m income last year

Low capital gains rate and US location of funds mean tax bill of £508,000 much less than under top income rate of 45%

Rishi Sunak in North Yorkshire, where he has a Grade-II listed manor with a private lake and heated swimming pool. Photograph: Ian Forsyth/PA

Rishi Sunak paid more than half a million pounds in tax in 2023 after making a £1.8m profit on his holding in a US investment fund, a summary of his tax affairs shows.

The prime minister published the document on Friday, showing he paid a tax bill of £508,308 in the financial year 2022-23 on overall earnings and gains of £2.23m.

This means he paid an effective tax rate of 23% in the UK - much lower than the top rate of 45% - because some income was taxed at source in the US and the rate of capital gains tax is lower at 20%.

Band	Taxable income	Tax rate
Personal Allowance	Up to £12,570	0%
Basic rate	£12,571 to £50,270	20%
Higher rate	£50,271 to £125,140	40%
Additional rate	over £125,140	45%

$(x-50271)*0.4+(50270-12570)*0.2 = 0.23*x$; $x = 73{,}932$

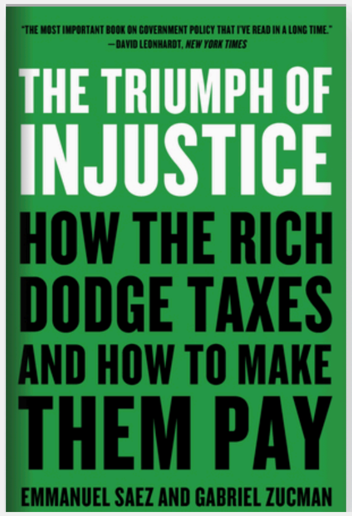

The Triumph of Injustice

How the Rich Dodge Taxes and How to Make Them Pay

by ***Emmanuel Saez*** *(Author, University of California, Berkeley),* ***Gabriel Zucman*** *(Author, University of California, Berkeley)*

"The most important book on government policy that I've read in a long time." —David Leonhardt, *New York Times*

https://taxjusticenow.org/

America's runaway inequality has an engine: our unjust tax system.

Even as they became fabulously wealthy, the ultra-rich have seen their taxes collapse to levels last seen in the 1920s.

Meanwhile, working-class Americans have been asked to pay more.

Make Your Own Tax Plan to see how adjusting taxes changes the progressivity of the tax system.

Dividend Policy: Why Pay Dividends?

- Historically and currently, this has been a very robust trend
  - Theoretically, in perfect markets dividend policies cannot affect value, and with higher tax dividends are almost always a bad deal compared with repurchases
- Many, many research papers look into this problem…
  - We combine annual stock market data for the most important equity markets of the last four centuries: the Netherlands and UK (1629–1812), UK (1813–1870), and US (1871–2015). We show that dividend yields are stationary and consistently forecast returns. (Golez and Koudijs, 2018)
  - Dividend payments remained prevalent even though repurchases were legal (Turner et al., 2013) and dividend taxation was present (Moortgat et al., 2017).

Panel A: Dividend-to-price ratio and five-year-ahead returns

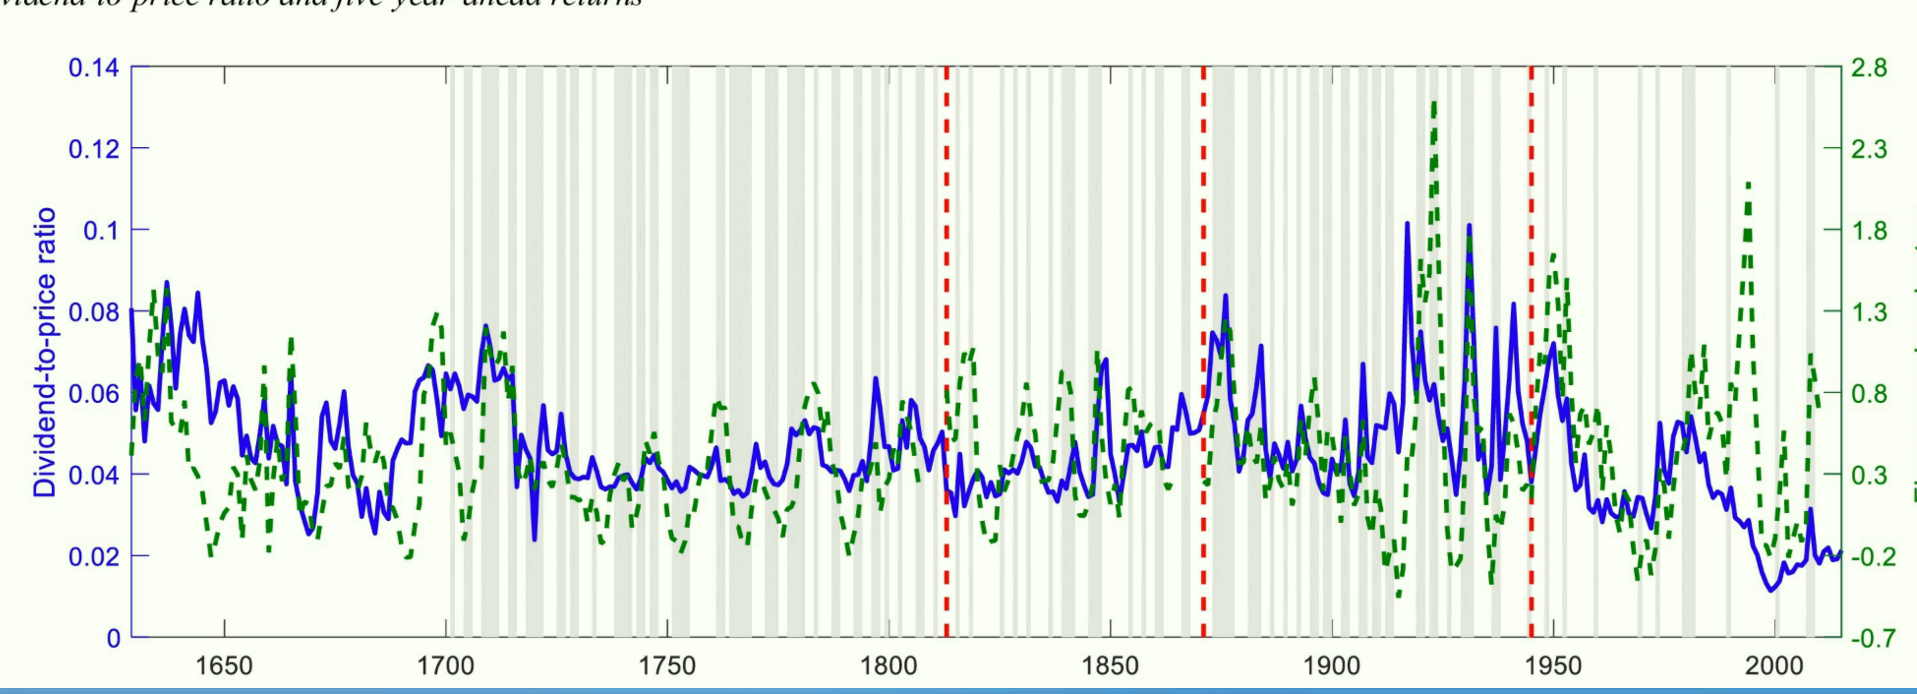

Repurchases

- A share repurchase is a transaction whereby a company buys back its own shares from the marketplace. A company might buy back its shares because management considers them undervalued.
- Buyback payments to investors may be tax-efficient if treated as capital gain/loss
- Ownership re-concentrated, shareholder alignment with management *may* be improved
- Used when companies have lots of cash, want to increase leverage.

Value the Markets

META Stock Soars After Dividend Announcement, Strong Earnings

Meta to reward investors, initiates stock buyback and dividend Plan. Explore Meta's strategic $50B share buyback and dividend plan,...

1. Suppose a firm operates for 2 periods (t=0,1) and has a 50% chance of each of two cash flows in period 1: £50 or £150. Investors are risk-neutral and the riskless interest rate is 10%. The firm is considering 2 financing strategies: all-equity and a zero-coupon bond with £60 face value.

a) Complete the following table

	All-equity	Leveraged
Payoff to shareholders in bad state	£50	£0
Payoff to shareholders in good state	£150	£90
Payoff to bondholders in bad state	£0	£50
Payoff to bondholders in good state	£0	£60
Value of debt	£0	½ * (50+60)/1.1 = £50
Value of equity	½ *(50+150)/1.1 = £90.91	½ *(0+90)/1.1 = £40.91
Value of firm	£90.91	£90.91

b) Now suppose investors are risk-averse. Another firm in the same line of business has a beta of 1.2 and the market rate of return is 15%. How would this modify the firm values computed in a? Need to use risk-adjusted discount rate $= r_f + \beta(r_{mkt} - r_f) = 16\%$. Using this instead of 10% in the above, the firm (independent of leverage) is now worth £86.21.

c) Now assume the firm's cash flows are *not* uncertain; instead, the firm generates £100/period forever. The riskless interest rate remains at 10%. The corporate tax rate (τ_C) is 30%, the personal tax rate on interest income (τ_i) is 40% and the tax rate on dividends is $\tau_e = 28\%$. Compute the firm value for each of the two strategies above (debt with face values of £0 and £60), and the value of the debt tax shield. After-tax cash flows per period are: $(1-\tau_e)(1-\tau_C)(£100 - r * D)$ to stockholders and $(1-\tau_i)r * D$ for bondholders. Adding these up, and using the fact that both income flows are certain and constant, the PDV is:

$$\left[\frac{(1-\tau_e)(1-\tau_C)}{(1-\tau_i)r}\right]£100 + \left[1 - \frac{(1-\tau_e)(1-\tau_C)}{(1-\tau_i)}\right]D$$

where the first (second) term is the unleveraged value (value of debt tax shield). The values are as follows

i. all equity = £840 + £0 = £840;
ii. leveraged = £789.90 + £60 = £849.60
iii. tax debt shield = £9.60

2. Imagine an economy with three groups of investors that face the following personal tax rates on interest payments:

Group:	Poor	Middle-class	Rich
Tax rate:	0%	40%	60%

Neither capital gains nor dividend payments on common equity are subject to personal tax. The corporate tax rate is 50%, but interest payments can be deducted from earnings for tax purposes. The economy has an infinite horizon, and annual cash flows (before interest and taxes) from a large number of identical corporations total £300 million. Each of the three investor groups has the same wealth, and regardless of capital structure they always invest the same amount in corporate securities. We assume they demand a minimum after-tax return of 10% on any security.

(a) Suppose all corporations are completely equity financed. What is the total tax bill? What is the combined firm value? How much is equity worth in total? Total tax = 50% * £300 million = £150 million; firm value = £150 million / 10% = £1.5 billion = value of equity.

(b) Suppose a firm decides to issue a certain amount of debt instead of equity, so that it now pays interest of £1 million each year, which investor group would hold this debt and why? What would be the interest rate? How would the value of the firm change? The 'poor' group is best placed to hold debt, because it pays no tax on interest. The interest will be that which makes them indifferent between debt and equity – namely 10%. The extra cash flow to the company would be £500,000/year, because taxable cash flow is now £1 million less, which lowers tax liability by 500,000. At 10% interest, this increases the firm's value by £5 million.

(c) Suppose one of the three groups holds all outstanding debt in the economy, but no equity. The firms pay bondholders interest at 10%. Which group holds the debt? What is total firm value? What are the total values of debt and equity? How much annual tax revenue does HMT receive? [Reminder: The total funds invested have to be the same for all three groups.] Again, it is the 'poor' who hold the debt – the other groups would net less than 10% after taxes, so the relevant interest tax rate is $\tau_i = 0\% = \tau_e$, and $\tau_C = 50\%$. Using the same formula as in 1c, but replacing £100 with cash flow CF, and taking account of the fact that all three groups invest the same amount (so $D = E/2$ and $V(D) = D + E = 3D$), we get $D = £1500/2.5 = £600$; equity value is £1200, corporate value is £1800, and tax revenue is £120

(d) Characterise qualitatively the (Miller) equilibrium in this economy. Starting from the situation in (c), which second group of investors can be induced to hold debt and how? Is it conceivable that all groups hold debt in equilibrium? The tax debt shield is not used up in c, so further trades are possible. Since each group will invest the same amount by assumption, the 'Middle-class' group will have to acquire debt. They will demand a pre-tax interest rate of 16.67% (to give post-tax returns of 10%). The equilibrium will have the poor owning only debt, the middle class a mix of debt and equity and the rich only equity. The rich would only hold debt if pre-tax interest were at least 25%, but the resulting cost of capital would be larger than the tax advantage.

Starting on Game Theory

- Define a problem
  - A finite set of players
    - Who have their own strategy space
  - A strategy space (also called action space)
    - Contains all possible strategies for each player
    - The vector of strategies for all players is a strategy profile
  - A payoff function (also called rewards function)
    - A mapping from a strategy profile to a real number
- Example

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Figure 2.18: The matching pennies

Cheap Talk Model

- Games with incomplete information: Perfect Bayesian Equilibrium
  - Both players choose their best responses
  - Their beliefs follow Bayes' rule where possible: $\mathbb{P}(B \mid A) \equiv \dfrac{\mathbb{P}(B \cap A)}{\mathbb{P}(A)}$
  - Another two components:
    - Θ_i is player i's type space and every element $\theta_i \in \Theta_i$ is a type
    - $v_i: A \times \Theta \to R$ is i's payoff function; where $A = A_1 \times \ldots \times A_n$ is the set of action profiles and $\Theta = \Theta_1 \times \ldots \times \Theta_n$ is the set of type profiles
- Player 1 has private information and the payoffs exhibit common values, so that both players' payoffs depend on player 1's private information. Player 1's action is a message that has no direct effect on payoffs.

1. Nature selects a type of player 1 $\theta \in \Theta$ from some common-knowledge distribution p.
2. Player 1 learns θ and chooses some message (action) $a_1 \in A_1$.
3. Player 2 observes message a_1 and chooses action $a_2 \in A_2$.
4. Payoffs $v_1(a_2, \theta)$ and $v_2(a_2, \theta)$ are realized.

A Motivating Example

To formalize the game we can think of me as player 1 , who is the sender of information, and my friend as player 2, who is the receiver of information, and we can imagine that expected traffic conditions are given by $\theta \in \{1,3,5\}$, where 1 is bad, 3 is mediocre, and 5 is good. Player 1 knows the true value of θ, but player 2 knows only the prior distribution of θ. Player 1 transmits a message (his action) to player 2 about the traffic conditions. Player 2 then chooses an action (where to live) $a_2 \in A_2 = \{1,2,3,4,5\}$, where 1 is San Francisco, 5 is Palo Alto, and 2, 3, and 4 are towns in between the two cities in that order.
The preferences of player 2 are described by the following payoff function: [1]

$$v_2(a_2, \theta) = 5 - (\theta - a_2)^2.$$

Notice that player 2 has a clear best response: given any level of traffic, he wants to choose his residence location equal to the traffic level. That is, his optimal choice is $a^*(\theta) = \theta$. To capture the fact that player 1 is biased toward having player 2 live closer to location 5 , the preferences of player 1 are given by the following payoff function:

$$v_1(a_2, \theta) = 5 - (\theta + b - a_2)^2,$$

where $b > 0$ is the bias of player 1 .

This is a dynamic game of incomplete information: player 1's type, or the state of the world θ, is known only to player 1 , while player 2 knows only the distribution of θ. Player 1's type affects both his payoff and the payoff of player 2, making this a common-values game. Player 1's action set includes messages that he can transmit to player 2, and player 2's action set includes choosing where to live. To further fix ideas, imagine that player 1 is restricted to sending only one of three messages corresponding to one of these states of nature: $a_1 \in A_1 = \{1,3,5\}$.[2]

A Motivating Example

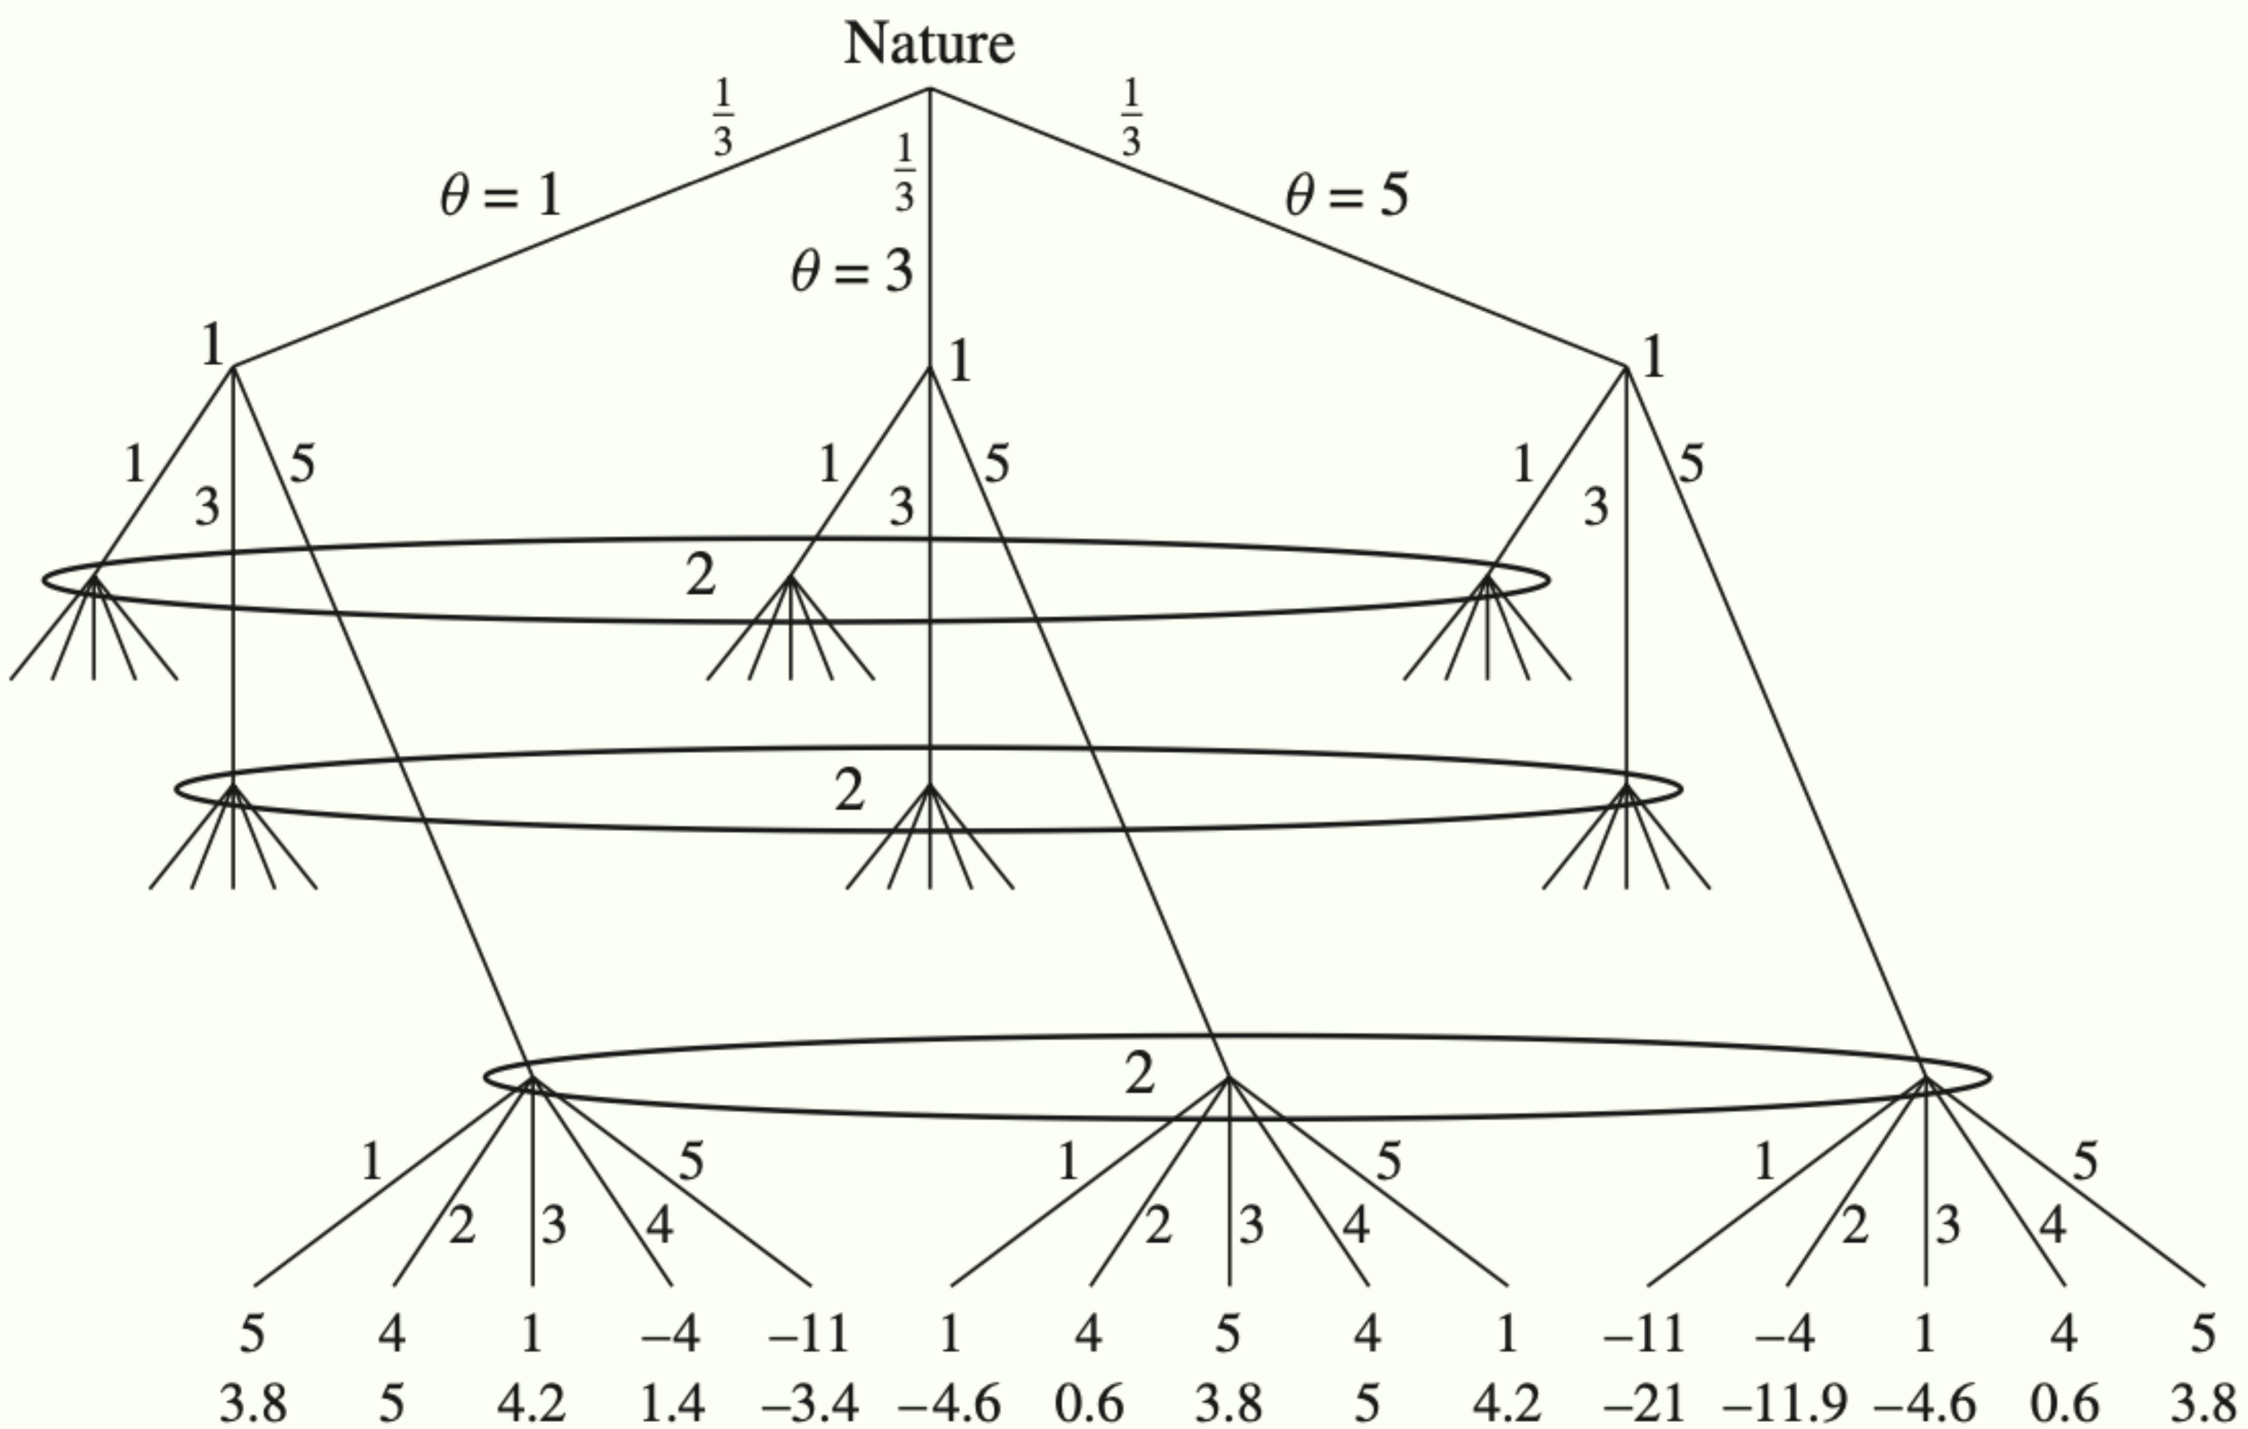

FIGURE 18.1 The commuting conditions information-transmission game.

A Motivating Example

Claim 18.1 *There is no perfect Bayesian equilibrium in which player 1 reports the true state of the world.*

Proof Assume in negation that player 1 truthfully reporting $a_1 = \theta$ is part of a perfect Bayesian equilibrium. It therefore must follow that when player 1 sends the message $a_1 \in \{1, 3, 5\}$ player 2 chooses $a_2 = a_1$. We saw that when $\theta = 3$ player 1 prefers $a_2 = 5$ over $a_2 = 3$. But if player 1 believes that player 2 will follow his advice then when $\theta = 3$ player 1's best response is $a_1 = 5$, a contradiction. ■

The intuition behind this result is simple, and easily generalizes to all such information-transmission games in which there is a bias between the sender and receiver with respect to the receiver's optimal choice. If it is indeed the case that the sender is reporting information truthfully, then it is in the receiver's best interest to take the sender's information at face value. But then if the receiver is acting in this way the sender has an incentive to lie.[5] The next observation is also quite straightforward:

A Motivating Example

Claim 18.2 *There exists a **babbling equilibrium** in which player 1's message reveals no information and player 2 chooses an action to maximize his expected utility given his prior belief.*

Proof To construct the babbling (perfect Bayesian) equilibrium let player 1's strategy be to send a message $a_1 \in \{1, 3, 5\}$ with equal probability of $\frac{1}{3}$ each regardless of θ. This means that the message is completely uninformative: player 2 knows that regardless of the message, $\Pr\{\theta\} = \frac{1}{3}$ for all $\theta \in \{1, 3, 5\}$. This implies that player 2 maximizes his expected payoff,

$$\max_{a_2 \in \{1,2,3,4,5\}} Ev_2(a_2, \theta) = 5 - \tfrac{1}{3}(-(1 - a_2)^2) + \tfrac{1}{3}(-(3 - a_2)^2) + \tfrac{1}{3}(-(5 - a_2)^2),$$

which is maximized when $a_2 = 3$.[6] Because each of player 2's information sets is reached with positive probability, player 2's beliefs are well defined by Bayes' rule everywhere, and player 1 cannot change these beliefs by changing his strategy.[7] Hence player 1 is indifferent between each of the three messages and is therefore playing a best response. ∎

Claims 18.1 and 18.2 paint a rather disappointing picture for our simple game. Not only will truthful messages never be part of an equilibrium (claim 18.1), but there is an equilibrium in which player 1's valuable private information will have no effect on player 2's choice. The remaining question is whether there are other equilibria in which there is *some* valuable information transmitted from the sender to the receiver.

A Motivating Example — Continuous Cheap Talk

The game is basically the same as the one described in the previous section, with the exception of $\theta \in \Theta = [0, 1]$ and the assumption that the state of the world θ is uniformly distributed on $[0, 1]$. Let player 2's action set include all real numbers, $a_2 \in \mathbb{R}$. Player 1's action set can be any arbitrary set of messages A_1, but it will be convenient to let $A_1 = [0, 1]$ so that the message space conforms with the state space Θ. Player 2's payoff is $v_2(a_2, \theta) = -(a_2 - \theta)^2$, while player 1's payoff is $v_1(a_2, \theta) = -(a_2 + b - \theta)^2$, which implies that for any given value of $\theta \in [0, 1]$, player 2's optimal choice is $a_2 = \theta$, while player 1's is $a_2 = \theta + b$. The payoff functions of the two players are depicted in Figure 18.2.

As in the finite example, both players would prefer a higher action to be taken when the state θ is higher, but player 1 has a constant bias, making him prefer even higher actions than player 2. This immediately implies that claim 18.1 generalizes to the continuous setting because of the same reasoning: If player 2 believes that player 1 is reporting θ truthfully, then player 2's best response is to choose $a_2 = \theta$. But if this is player 2's strategy then player 1 will report $a_1 = \theta + b$ for any $b \neq 0$. Hence there can never be a fully truthful equilibrium. Not surprisingly a babbling equilibrium still exists:

Claim 18.4 *There exists a babbling perfect Bayesian equilibrium in which player 1's message reveals no information and player 2 chooses an action to maximize his expected utility given his prior belief.*

Proof We construct the equilibrium in a similar way to the finite case. Let player 1's strategy be to send a message $a_1 = a_1^B \in [0, 1]$ regardless of θ. This means that the message is completely uninformative and player 2 believes that θ is distributed uniformly on $[0, 1]$. This implies that, conditional on receiving the message a_1^B, player 2 maximizes his expected payoff,

$$\max_{a_2 \in \mathbb{R}} Ev_2(a_2, \theta) = \int_0^1 -(\theta - a_2)^2 d\theta \ = -\tfrac{1}{3} + 2a_2 - a_2^2,$$

which is maximized when $a_2 = \frac{1}{2}$. Let player 2's off-equilibrium-path beliefs be $\Pr\{\theta = \frac{1}{2} | a_1 \neq a_1^B\} = 1$ so that his off-the-equilibrium-path best response to any other message is $a_2 = \frac{1}{2}$ as well. It is easy to see that player 1 is indifferent between any of his messages and hence choosing $a_1 = a_1^B$ is a best response. ■

We see that the continuous-space cheap-talk model has the same two extreme results demonstrated for the discrete-space game: there is no truthful equilibrium and there is always a babbling equilibrium. The question then is, how much information can the sender, player 1, credibly transmit to the receiver, player 2? We begin by constructing a perfect Bayesian equilibrium in which player 1 uses one of two messages, a_1' and a_1'', and player 2 chooses a different action following each message, $a_2(a_1') < a_2(a_1'')$.[8]

Claim 18.5 *In a two-message equilibrium player 1 must use a threshold strategy as follows: if $0 \le \theta \le \theta^*$ he chooses a_1', whereas if $\theta^* \le \theta \le 1$ he chooses a_1''.*

Proof For any θ player 1's payoffs from a_1' and a_1'' are as follows:

$$v_1(a_2(a_1'), \theta) = -(a_2(a_1') + b - \theta)^2$$

$$v_1(a_2(a_1''), \theta) = -(a_2(a_1'') + b - \theta)^2,$$

which implies that the extra gain from choosing a_1'' over a_1' is equal to

$$\Delta v_1(\theta) = -(a_2(a_1'') + b - \theta)^2 + (a_2(a_1') + b - \theta)^2.$$

The derivative of $\Delta v_1(\theta)$ is equal to $2(a_2(a_1'') - a_2(a_1')) > 0$ because $a_2(a_1') < a_2(a_1'')$. This implies that if type θ prefers to send message a_1'' over a_1' then every type $\theta' > \theta$ will also prefer a_1''. Similarly if type θ prefers to send message a_1' over a_1'' then so will every type $\theta' < \theta$. This in turn implies that if two messages are sent in equilibrium then there must be some threshold type θ^* as defined in claim 18.5. It follows that when $\theta = \theta^*$ player 1 must be indifferent between sending the two messages. ■

Claim 18.6 *In any two-message perfect Bayesian equilibrium in which player 1 is using a threshold θ^* strategy as described in claim 18.5, player 2's equilibrium best response is $a_2(a_1') = \frac{\theta^*}{2}$ and $a_2(a_1'') = \frac{1-\theta^*}{2}$.*

Proof This follows from player 2's posterior belief and from him playing a best response. In equilibrium player 2's posterior following a message a_1' is that θ is uniformly distributed on the interval $[0, \theta^*]$, and his posterior following a message a_1'' is that θ is uniformly distributed on the interval $[\theta^*, 1]$. Player 2 plays a best response if and only if he sets $a_2(a_1) = E[\theta | a_1]$, which proves the result. ■

Claim 18.7 *A two-message perfect Bayesian equilibrium exists if and only if $b < \frac{1}{4}$.*

Proof From claim 18.5 we know that when $\theta = \theta^*$ player 1 must be indifferent between his two messages so that

$$v_1(a_2(a_1'), \theta^*) = v_1(a_2(a_1''), \theta^*),$$

which from claim 18.6 and from the fact that $\frac{\theta^*}{2} < \theta^* < \frac{1-\theta^*}{2}$ is equivalent to

$$\theta^* + b - \frac{\theta^*}{2} = -\left(\theta^* + b - \frac{1-\theta^*}{2}\right). \tag{18.1}$$

The solution to (18.1) is $\theta^* = \frac{1}{4} - b$, which can result in a positive value of θ^* only if $b < \frac{1}{4}$. To complete the specification of off-the-equilibrium-path beliefs, let player 2's beliefs be $\Pr\{\theta = \frac{\theta^*}{2} | a_1 \notin \{a_1', a_1''\}\} = 1$, so that he chooses $a_2 = \frac{\theta^*}{2}$, which causes player 1 to be indifferent between sending the message a_1' and any other message $a_1 \notin \{a_1', a_1''\}$, implying that his threshold strategy is a best response. ■

Example from 2021 Exam

A government procurement officer is trying to decide how many doses of a new coronavirus vaccine to order. This decision will depend on the effectiveness of the vaccine, which will be determined by clinical trials conducted by a scientific advisor. You may assume that the effectiveness of the vaccine is given by a random variable ε, uniformly distributed on the interval $[E_0, E_0 + 1]$. The scientific advisor believes the utility of ordering a quantity Q is $U_A(Q|\varepsilon) = 1 + \varepsilon Q - Q^2$; if perfectly-informed about effectiveness, the procurement officer would value Q at $U_G(Q|\varepsilon) = 1 + (\varepsilon + \beta)Q - Q^2$ where β is a non-negative constant. After the trials, the government officer asks the scientific advisor to report on the vaccine's effectiveness and purchases the quantity that maximises his expected utility. The scientific advisor is not paid for his efforts.

a. How would you set up this problem? Can the government advisor be sure of purchasing the optimal quantity (according to his preferences)? If so, how? If not, why not? How does your answer depend on the size of β? (**15 marks**)

b. Suppose that the minimum effectiveness is $E_0 = 25\%$ and that $\beta = 5\%$. Find the 'babbing equilibrium' for this situation – how much will the government order and what expected utilities will the two parties get? (**7 marks**)

c. Now construct a two-part equilibrium – depending on the advice they receive, the government will place either a small order Q^S or a large order Q^L. At what reported level of effectiveness will the government switch its order size, and what are the values of Q^S and Q^L? (**10 marks**)

d. How would you find the most efficient equilibrium (you do not have to compute it explicitly, but should say how it could be identified)? (**10 marks**) a

This is a cheap talk problem; should note that first-best can be achieved only if $\beta = 0$. They should note that there is always an equilibrium in which the government ignores the advisor and purchases the a priori optimal amount $Q_0(E_0, E_0 + 1)$, which they compute in the next part. The optimal strategy is to partition the range of effectiveness into intervals $[x, y]$ and associate to each interval the order that maximises expected utility $Q^*(x,y) = \underset{Q}{\operatorname{argmax}} \int_x^y U_G(Q|\varepsilon)d\varepsilon$. The more intervals, the more efficient is the outcome, but the number (and thus the efficiency) are bounded above by a decreasing function of β. Finally, they should note that for any two adjacent intervals $[x, y]$ and $[y, z]$, the scientific adviser would be indifferent between the purchase levels for both intervals if she was convinced that the true effectiveness was exactly y – in other words $U_A(Q^*(x,y)|y) = U_A(Q^*(y,z)|y)$.

b

In this case, there is only one purchase level regardless of report. If the government believes that the true state is uniformly distributed on $[a, b]$, it's expected utility for purchasing Q is

$$EU_G(Q) = 1 + \left(\frac{\int_a^b \varepsilon d\varepsilon}{b-a} + \beta\right)Q - Q^2 = 1 + \left(\frac{a+b+2\beta}{2}\right)Q - Q^2$$

Optimal Q is $\frac{a+b+2\beta}{4}$. In this case, $Q = \frac{.25+1.25+.1}{4} = 0.4$, $U_G = 1.16$, $U_A = 1.14$.

c

Denote the critical report level by ε^*. The two order sizes are

$$Q^S = \frac{.25 + \varepsilon^* + 2 * .05}{4} = \frac{.35 + \varepsilon^*}{4}$$

$$Q^L = \frac{1.25 + \varepsilon^* + 2 * .05}{4} = \frac{1.35 + \varepsilon^*}{4}$$

ε^* is defined by the condition that the advisor should be indifferent between Q^S and Q^L when the true state is ε^*. Solving $U_A\left(\frac{.35+\varepsilon^*}{4}\middle|\varepsilon^*\right) = U_A\left(\frac{1.35+\varepsilon^*}{4}\middle|\varepsilon^*\right)$ for ε^* gives $\varepsilon =$ (in general, for any level of β), $\varepsilon^* = 0.75 + 2\beta$; in this case,

$$\varepsilon^* = 0.85$$

$$Q^S = 0.3$$

$$Q^L = 0.55$$

d

The most efficient equilibrium is the one with the greatest number of intervals, so they should look for the largest n s.t. there exists a sequence $0.25 = \varepsilon^1, \dots, \varepsilon^n = 1.25$ (or .26 for the 1% case) where

$$Q^i = \frac{\varepsilon^i + \varepsilon^{i+1} + 0.1}{4}$$

And for each $i = 1, \dots, n-1$

$$1 + \varepsilon^{i+1}\left[\frac{\varepsilon^i + \varepsilon^{i+1} + 0.1}{4}\right] - \left[\frac{\varepsilon^i + \varepsilon^{i+1} + 0.1}{4}\right]^2 = 1 + \varepsilon^{i+1}\left[\frac{\varepsilon^{i+1} + \varepsilon^{i+2} + 0.1}{4}\right] - \left[\frac{\varepsilon^{i+1} + \varepsilon^{i+2} + 0.1}{4}\right]^2$$