How to Make the Right Decisions without Knowing People's Preferences: An Introduction to Mechanism Design

E. Maskin

The Directors Roundtable
Shearman & Sterling Conference Center
New York
October, 2014

Theory of Mechanism Design –

Theory of Mechanism Design –

“engineering” part of economic theory

Theory of Mechanism Design –

"engineering" part of economic theory

- much of economic theory devoted to:

Theory of Mechanism Design –

"engineering" part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions

Theory of Mechanism Design –

“engineering” part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions
  - explaining/predicting outcomes that institutions generate

Theory of Mechanism Design –

“engineering” part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions
  - explaining/predicting outcomes that institutions generate
  - positive, predictive

Theory of Mechanism Design –

“engineering” part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions
  - explaining/predicting outcomes that institutions generate
  - positive, predictive

- mechanism design – reverses the direction

Theory of Mechanism Design –

“engineering” part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions
  - explaining/predicting outcomes that institutions generate
  - positive, predictive

- mechanism design – reverses the direction
  - begins by identifying desired outcomes (goals)

Theory of Mechanism Design –

“engineering” part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions
  - explaining/predicting outcomes that institutions generate
  - positive, predictive

- mechanism design – reverses the direction
  - begins by identifying desired outcomes (goals)
  - asks whether institutions (mechanisms) could be designed to achieve goals

Theory of Mechanism Design –

“engineering” part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions
  - explaining/predicting outcomes that institutions generate
  - positive, predictive

- mechanism design – reverses the direction
  - begins by identifying desired outcomes (goals)
  - asks whether institutions (mechanisms) could be designed to achieve goals
  - if so, what forms would institutions take?

Theory of Mechanism Design –

“engineering” part of economic theory

- much of economic theory devoted to:
  - understanding existing economic institutions
  - explaining/predicting outcomes that institutions generate
  - positive, predictive

- mechanism design – reverses the direction
  - begins by identifying desired outcomes (goals)
  - asks whether institutions (mechanisms) could be designed to achieve goals
  - if so, what forms would institutions take?
  - normative, prescriptive

For example, suppose

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob
- goal: divide so that each child is happy with his/her portion

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob
- goal: divide so that each child is happy with his/her portion
  - Bob thinks he has got at least half

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob
- goal: divide so that each child is happy with his/her portion
  - Bob thinks he has got at least half
  - Alice thinks she has got at least half

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob
- goal: divide so that each child is happy with his/her portion
  - Bob thinks he has got at least half
  - Alice thinks she has got at least half

  call this *fair division*

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob
- goal: divide so that each child is happy with his/her portion
  - Bob thinks he has got at least half
  - Alice thinks she has got at least half

  call this *fair division*

- If mother knows that the kids see the cake in same way she does, simple solution:

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob
- goal: divide so that each child is happy with his/her portion
  - Bob thinks he has got at least half
  - Alice thinks she has got at least half

  call this *fair division*

- If mother knows that the kids see the cake in same way she does, simple solution:
  - she divides equally (in her view)

For example, suppose

- mother wants to divide cake between 2 children, Alice and Bob
- goal: divide so that each child is happy with his/her portion
  - Bob thinks he has got at least half
  - Alice thinks she has got at least half

  call this *fair division*

- If mother knows that the kids see the cake in same way she does, simple solution:
  - she divides equally (in her view)
  - gives each kid a portion

- But what if, say, Bob sees cake differently from mother?

- But what if, say, Bob sees cake differently from mother?
  - *she* thinks she’s divided it equally

- But what if, say, Bob sees cake differently from mother?
  - *she* thinks she’s divided it equally
  - but *he* thinks piece he’s received is smaller than Alice’s

- But what if, say, Bob sees cake differently from mother?
    - *she* thinks she’s divided it equally
    - but *he* thinks piece he’s received is smaller than Alice’s

- difficulty: mother wants to achieve *fair division*

- But what if, say, Bob sees cake differently from mother?
  - *she* thinks she’s divided it equally
  - but *he* thinks piece he’s received is smaller than Alice’s

- difficulty: mother wants to achieve *fair division*
  - but doesn’t have enough information to do this on her own

- But what if, say, Bob sees cake differently from mother?
    - *she* thinks she’s divided it equally
    - but *he* thinks piece he’s received is smaller than Alice’s

- difficulty: mother wants to achieve *fair division*
    - but doesn’t have enough information to do this on her own
    - in effect, doesn’t know which division is fair

- Can she design a mechanism (procedure) for which outcome will be a fair division?

- Can she design a mechanism (procedure) for which outcome will be a fair division?

(even though she doesn’t know what is fair herself ?)

- Can she design a mechanism (procedure) for which outcome will be a fair division?

  (even though she doesn’t know what is fair herself ?)

- Age-old problem

- Can she design a mechanism (procedure) for which outcome will be a fair division?

  (even though she doesn’t know what is fair herself ?)

- Age-old problem
  - Lot and Abraham dividing grazing land

Age-old solution:

Age-old solution:

  - have *Bob* divide the cake in two

Age-old solution:

- have *Bob* divide the cake in two
- have *Alice* choose one of the pieces

Age-old solution:

- have *Bob* divide the cake in two
- have *Alice* choose one of the pieces

Why does this work?

Age-old solution:

  - have *Bob* divide the cake in two
  - have *Alice* choose one of the pieces

Why does this work?

- Bob will divide so that pieces are equal in his eyes

Age-old solution:

  - have *Bob* divide the cake in two
  - have *Alice* choose one of the pieces

Why does this work?

- Bob will divide so that pieces are equal in his eyes
  - if one of the pieces were bigger, then Alice would take that one

Age-old solution:

  - have *Bob* divide the cake in two
  - have *Alice* choose one of the pieces

Why does this work?

- Bob will divide so that pieces are equal in his eyes
  - if one of the pieces were bigger, then Alice would take that one
- So whichever piece Alice takes, Bob will be happy with other

Age-old solution:

  - have *Bob* divide the cake in two
  - have *Alice* choose one of the pieces

Why does this work?

- Bob will divide so that pieces are equal in his eyes
  - if one of the pieces were bigger, then Alice would take that one
- So whichever piece Alice takes, Bob will be happy with other
- And Alice will be happy with her own choice because if she thinks pieces unequal, can take bigger one

Example illustrates key features of mechanism design:

Example illustrates key features of mechanism design:

- mechanism designer herself *doesn’t know* in advance what outcomes are optimal

Example illustrates key features of mechanism design:

- mechanism designer herself *doesn’t know* in advance what outcomes are optimal
- so must proceed indirectly through a mechanism

Example illustrates key features of mechanism design:

- mechanism designer herself *doesn’t know* in advance what outcomes are optimal
- so must proceed indirectly through a mechanism
  - have participants *themselves* generate information needed to identify optimal outcome

Example illustrates key features of mechanism design:

- mechanism designer herself *doesn't know* in advance what outcomes are optimal
- so must proceed indirectly through a mechanism
  - have participants *themselves* generate information needed to identify optimal outcome
- complication: participants don't care about *mechanism designer's goals*

Example illustrates key features of mechanism design:

- mechanism designer herself *doesn’t know* in advance what outcomes are optimal
- so must proceed indirectly through a mechanism
  - have participants *themselves* generate information needed to identify optimal outcome
- complication: participants don’t care about *mechanism designer’s goals*
  - have their own *objectives*

Example illustrates key features of mechanism design:

- mechanism designer herself *doesn’t know* in advance what outcomes are optimal
- so must proceed indirectly through a mechanism
  - have participants *themselves* generate information needed to identify optimal outcome
- complication: participants don’t care about *mechanism designer’s goals*
  - have their own *objectives*
- so mechanism must be *incentive compatible*

Example illustrates key features of mechanism design:

- mechanism designer herself *doesn't know* in advance what outcomes are optimal
- so must proceed indirectly through a mechanism
  - have participants *themselves* generate information needed to identify optimal outcome
- complication: participants don't care about *mechanism designer's goals*
  - have their own *objectives*
- so mechanism must be *incentive compatible*
  - must reconcile social and individual goals

Second Example:

Second Example:

Suppose government wants to sell right (license) to transmit on band of radio frequencies

Second Example:

Suppose government wants to sell right (license) to transmit on band of radio frequencies

(real-life issue for many governments)

Second Example:

Suppose government wants to sell right (license) to transmit on band of radio frequencies

(real-life issue for many governments)

- several telecommunication companies interested in license

Second Example:

Suppose government wants to sell right (license) to transmit on band of radio frequencies

(real-life issue for many governments)

- several telecommunication companies interested in license
- goal of government: to put transmitting license in hands of company that values it most (“efficient” outcome)

Second Example:

Suppose government wants to sell right (license) to transmit on band of radio frequencies

(real-life issue for many governments)

- several telecommunication companies interested in license
- goal of government: to put transmitting license in hands of company that values it most (“efficient” outcome)
- but government doesn’t know how much each company values it (so doesn’t know best outcome)

Government could ask each company how much it values license

Government could ask each company how much it values license

- but if company thinks its chances of getting license go up when it states higher value, has incentive to *exaggerate* value

Government could ask each company how much it values license

- but if company thinks its chances of getting license go up when it states higher value, has incentive to *exaggerate* value
- so no guarantee of identifying company that values it most

- government could have

- government could have
  - each company make a bid for license

- government could have
  - each company make a bid for license
  - high bidder wins license

- government could have
  - each company make a bid for license
  - high bidder wins license
  - winner pays bid

- government could have
    - each company make a bid for license
    - high bidder wins license
    - winner pays bid
- but this mechanism won’t work either

- government could have
  - each company make a bid for license
  - high bidder wins license
  - winner pays bid
- but this mechanism won’t work either
  - companies have incentive to *understate*

- government could have
  - each company make a bid for license
  - high bidder wins license
  - winner pays bid
- but this mechanism won’t work either
  - companies have incentive to *understate*
- suppose license worth $10m to Telemax, then

- government could have
  - each company make a bid for license
  - high bidder wins license
  - winner pays bid
- but this mechanism won’t work either
  - companies have incentive to *understate*
- suppose license worth $10m to Telemax, then
  - if Telemax bids $10m and wins, gets
    $10m − $10m = 0

- government could have
  - each company make a bid for license
  - high bidder wins license
  - winner pays bid
- but this mechanism won’t work either
  - companies have incentive to *understate*
- suppose license worth \$10m to Telemax, then
  - if Telemax bids \$10m and wins, gets
    \$10m − \$10m = 0
- so Telemax will bid *less* than \$10m

- government could have
  - each company make a bid for license
  - high bidder wins license
  - winner pays bid
- but this mechanism won’t work either
  - companies have incentive to *understate*
- suppose license worth $10m to Telemax, then
  - if Telemax bids $10m and wins, gets $10m − $10m = 0
- so Telemax will bid *less* than $10m
- but if all bidders are understating, no guarantee that winner will be company that values license most

Solution:

Solution:

- every company makes bid for license

Solution:

- every company makes bid for license
- winner is high bidder

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate
  - doesn’t pay bid anyway

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate
  - doesn’t pay bid anyway
  - if understates, may lose license

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate
  - doesn’t pay bid anyway
  - if understates, may lose license
- Has no incentive to overstate

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate
  - doesn’t pay bid anyway
  - if understates, may lose license
- Has no incentive to overstate
  - If bids $12m, will now win if other company bids $11m

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate
  - doesn’t pay bid anyway
  - if understates, may lose license
- Has no incentive to overstate
  - If bids $12m, will now win if other company bids $11m
  - But overpays

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate
  - doesn’t pay bid anyway
  - if understates, may lose license
- Has no incentive to overstate
  - If bids $12m, will now win if other company bids $11m
  - But overpays
- So best to bid *exactly* what license worth

Solution:

- every company makes bid for license
- winner is high bidder
- winner pays *second-highest* bid
  - so if 3 bidders and bids are
    $10m, $8m, and $5m,
    winner is company that bids $10m
  - but pays only $8m
- Now company has no incentive to understate
  - doesn’t pay bid anyway
  - if understates, may lose license
- Has no incentive to overstate
  - If bids $12m, will now win if other company bids $11m
  - But overpays
- So best to bid *exactly* what license worth
- And winner will be company that values license most

Final Example

Final Example

Consider society with

Final Example

Consider society with

- 2 consumers of energy – Alice and Bob

Final Example

Consider society with

- 2 consumers of energy – Alice and Bob
- Energy authority – must choose public energy source

Final Example

Consider society with

- 2 consumers of energy – Alice and Bob

- Energy authority – must choose public energy source
    - gas

Final Example

Consider society with

- 2 consumers of energy – Alice and Bob
- Energy authority – must choose public energy source
  - gas
  - oil

Final Example

Consider society with

- 2 consumers of energy – Alice and Bob

- Energy authority – must choose public energy source
  - gas
  - oil
  - nuclear power

Final Example

Consider society with

- 2 consumers of energy – Alice and Bob
- Energy authority – must choose public energy source
  - gas
  - oil
  - nuclear power
  - coal

Two states of world

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

Two states of world

state 1   consumers weight future lightly (future relatively unimportant)

state 2   consumers weight future heavily (future relatively important)

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1: favors gas over oil, oil over coal, and coal over nuclear

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1: favors gas over oil, oil over coal, and coal over nuclear

In state 2: favors nuclear over gas, gas over coal, and coal over oil

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1: favors gas over oil, oil over coal, and coal over nuclear

In state 2: favors nuclear over gas, gas over coal, and coal over oil

− technical advances expected to make gas, coal, and especially nuclear easier to use in future compared with oil

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1: favors gas over oil, oil over coal, and coal over nuclear

In state 2: favors nuclear over gas, gas over coal, and coal over oil

− technical advances expected to make gas, coal, and especially nuclear easier to use in future compared with oil

Bob – cares more about safety

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1: favors gas over oil, oil over coal, and coal over nuclear

In state 2: favors nuclear over gas, gas over coal, and coal over oil

− technical advances expected to make gas, coal, and especially nuclear easier to use in future compared with oil

Bob – cares more about safety

In state 1: favors nuclear over oil, oil over coal, and coal over gas

Two states of world

state 1   consumers weight future lightly (future relatively unimportant)

state 2   consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1:   favors gas over oil, oil over coal, and coal over nuclear

In state 2:   favors nuclear over gas, gas over coal, and coal over oil

− technical advances expected to make gas, coal, and especially nuclear easier to use in future compared with oil

Bob – cares more about safety

In state 1:   favors nuclear over oil, oil over coal, and coal over gas

In state 2:   favors oil over gas, gas over coal, and coal over nuclear

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1: favors gas over oil, oil over coal, and coal over nuclear

In state 2: favors nuclear over gas, gas over coal, and coal over oil

− technical advances expected to make gas, coal, and especially nuclear easier to use in future compared with oil

Bob – cares more about safety

In state 1: favors nuclear over oil, oil over coal, and coal over gas

In state 2: favors oil over gas, gas over coal, and coal over nuclear

− disposal of nuclear waste will loom large

Two states of world

state 1 consumers weight future lightly (future relatively unimportant)

state 2 consumers weight future heavily (future relatively important)

Alice – cares mainly about convenience

In state 1: favors gas over oil, oil over coal, and coal over nuclear

In state 2: favors nuclear over gas, gas over coal, and coal over oil

− technical advances expected to make gas, coal, and especially nuclear easier to use in future compared with oil

Bob – cares more about safety

In state 1: favors nuclear over oil, oil over coal, and coal over gas

In state 2: favors oil over gas, gas over coal, and coal over nuclear

− disposal of nuclear waste will loom large

− gas will become safer

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

- energy authority

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

- energy authority
  - wants source that makes good compromise between consumers' views

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

- energy authority
    - wants source that makes good compromise between consumers’ views
    - so, oil is social optimum in state 1

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

- energy authority
  - wants source that makes good compromise between consumers’ views
  - so, oil is social optimum in state 1
  - gas is social optimum in state 2

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

- energy authority
  - wants source that makes good compromise between consumers' views
  - so, oil is social optimum in state 1
  - gas is social optimum in state 2
- but suppose authority *does not know* state

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

- energy authority
  - wants source that makes good compromise between consumers' views
  - so, oil is social optimum in state 1
  - gas is social optimum in state 2
- but suppose authority *does not know* state
  - then doesn't know whether oil or gas better

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

oil optimal

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

gas optimal

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear
oil optimal		gas optimal	

- authority could ask Alice or Bob about state

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

oil optimal

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

gas optimal

- authority could ask Alice or Bob about state
  - but Alice has incentive to say “state 2” *regardless* of truth

    always prefers gas to oil

    gas optimal in state 2

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear
oil optimal		gas optimal	

- authority could ask Alice or Bob about state
  - but Alice has incentive to say “state 2” *regardless* of truth

    always prefers gas to oil

    gas optimal in state 2
  - Bob always has incentive to say “state 1”

    always prefers oil to gas

    oil optimal state 1

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

oil optimal

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

gas optimal

– authority could ask Alice or Bob about state
  - but Alice has incentive to say “state 2” *regardless* of truth

    always prefers gas to oil

    gas optimal in state 2
  - Bob always has incentive to say “state 1”

    always prefers oil to gas

    oil optimal state 1

So, simply asking consumers to reveal actual state too naive a mechanism

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Authority can have consumers participate in the mechanism given by table

	Bob	
Alice	oil	coal
	nuclear	gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Authority can have consumers participate in the mechanism given by table

	Bob	
Alice	oil	coal
	nuclear	gas

- Alice – can choose top row or bottom row

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Authority can have consumers participate in the mechanism given by table

Bob

Alice

oil	coal
nuclear	gas

- Alice – can choose top row or bottom row
- Bob – can choose left column or right column

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear
social optimum: oil		social optimum: gas	

Authority can have consumers participate in the mechanism given by table

Bob

Alice	oil	coal
	nuclear	gas

- Alice – can choose top row or bottom row
- Bob – can choose left column or right column
- outcomes given by table entries

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Authority can have consumers participate in the mechanism given by table

	Bob	
Alice	oil	coal
	nuclear	gas

- Alice – can choose top row or bottom row
- Bob – can choose left column or right column
- outcomes given by table entries
- If state 1 holds
  - Alice will prefer top row if Bob plays left column
  - Bob will always prefer left column
  - so (Alice plays top, Bob plays left) is Nash equilibrium
    - neither participant has incentive to change unilaterally to another strategy
  - In fact, it is *unique* Nash equilibrium

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Authority can have consumers participate in the mechanism given by table

		Bob
Alice	oil	coal
	nuclear	gas

- Alice – can choose top row or bottom row
- Bob – can choose left column or right column
- outcomes given by table entries
- If state 1 holds
  Alice will prefer top row if Bob plays left column
  Bob will always prefer left column
  so (Alice plays top, Bob plays left) is Nash equilibrium
    neither participant has incentive to change unilaterally to another strategy
  In fact, it is *unique* Nash equilibrium
    - so good prediction of what Alice and Bob will do

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Bob

Alice	oil	coal
	nuclear	gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Bob

Alice

oil	coal
nuclear	gas

So, in state 1:

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

So, in state 1:

- expect that

  Alice will play top strategy

  Bob will play left strategy

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

So, in state 1:

- expect that
  Alice will play top strategy
  Bob will play left strategy
- outcome is oil

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

So, in state 1:

- expect that
  Alice will play top strategy
  Bob will play left strategy
- outcome is oil
- oil is social optimum

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

Similarly, in state 2:

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Bob

Alice	oil	coal
	nuclear	gas

Similarly, in state 2:

- expect that

  Alice will play bottom strategy

  Bob will play right strategy

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Bob

Alice

oil	coal
nuclear	gas

Similarly, in state 2:

- expect that
  Alice will play bottom strategy
  Bob will play right strategy
- outcome is gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Bob

Alice	oil	coal
	nuclear	gas

Similarly, in state 2:

- expect that
  Alice will play bottom strategy
  Bob will play right strategy
- outcome is gas
- gas is social optimum

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

- Thus, in *either state*, mechanism achieves social optimum, even though

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Bob

Alice	oil	coal
	nuclear	gas

- Thus, in *either state*, mechanism achieves social optimum, even though
  - mechanism designer doesn’t know the state herself

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

Bob

Alice	oil	coal
	nuclear	gas

- Thus, in *either state*, mechanism achieves social optimum, even though
  - mechanism designer doesn’t know the state herself
  - Alice and Bob interested in own ends (not social goal)

State 1

Alice	Bob
gas	nuclear
oil	oil
coal	coal
nuclear	gas

social optimum: oil

State 2

Alice	Bob
nuclear	oil
gas	gas
coal	coal
oil	nuclear

social optimum: gas

	Bob	
Alice	oil	coal
	nuclear	gas

- Thus, in *either state*, mechanism achieves social optimum, even though
  - mechanism designer doesn’t know the state herself
  - Alice and Bob interested in own ends (not social goal)
- We say that mechanism *implements* the designer’s goals (oil in state 1, gas in state 2)

- Have shown you mechanisms in the cake and telecommunication examples

- Have shown you mechanisms in the cake and telecommunication examples
- But analysis may seem a bit *ad hoc*

- Have shown you mechanisms in the cake and telecommunication examples
- But analysis may seem a bit *ad hoc*
- Examples prompt questions:

- Have shown you mechanisms in the cake and telecommunication examples
- But analysis may seem a bit *ad hoc*
- Examples prompt questions:
  - is there a *general* way of determining whether or not a given goal is implementable?

- Have shown you mechanisms in the cake and telecommunication examples
- But analysis may seem a bit *ad hoc*
- Examples prompt questions:
  - is there a *general* way of determining whether or not a given goal is implementable?
  - if it *is* implementable, can we find a mechanism that implements it?

- Have shown you mechanisms in the cake and telecommunication examples
- But analysis may seem a bit *ad hoc*
- Examples prompt questions:
  - is there a *general* way of determining whether or not a given goal is implementable?
  - if it *is* implementable, can we find a mechanism that implements it?
- Answer: yes to both questions

  see Maskin “Nash Equilibrium and Welfare Optimality,” 1977

- Have looked at 2 applications of mechanism design theory

- Have looked at 2 applications of mechanism design theory
- Many other potential applications

- Have looked at 2 applications of mechanism design theory
- Many other potential applications

1) International treaty on greenhouse gas emissions

- Have looked at 2 applications of mechanism design theory
- Many other potential applications

1) International treaty on greenhouse gas emissions
2) Design of presidential elections

American Economic Review 2008, 98:3, 567–576
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.1.567

Mechanism Design: How to Implement Social Goals[†]

By Eric S. Maskin*

The theory of mechanism design can be thought of as the "engineering" side of economic theory. Much theoretical work, of course, focuses on *existing* economic institutions. The theorist wants to explain or forecast the economic or social outcomes that these institutions generate. But in mechanism design theory the direction of inquiry is reversed. We begin by identifying our desired outcome or *social goal*. We then ask whether or not an appropriate institution (mechanism) could be designed to attain that goal. If the answer is yes, then we want to know what form that mechanism might take.

In this paper, I offer a brief introduction to the part of mechanism design called *implementation theory*, which, given a social goal, characterizes when we can design a mechanism whose *predicted* outcomes (i.e., the set of equilibrium outcomes) coincide with the *desirable* outcomes, according to that goal. I try to keep technicalities to a minimum, and usually confine them to footnotes.[1]

I. Outcomes, Goals, and Mechanisms

What we mean by an "outcome" will naturally depend on the context. Thus, for a government charged with delivering public goods, an outcome will consist of the quantities provided of such goods as intercity highways, national defense and security, environmental protection, and public education, together with the arrangements by which they are financed. For an electorate seeking to fill a political office, an outcome is simply the choice of a candidate for that office. For an auctioneer selling a collection of assets, an outcome corresponds to an allocation of these assets across potential buyers, together with the payments that these buyers make. Finally, in the case of a home buyer and a builder contemplating the construction of a new house, an outcome is a specification of the house's characteristics and the builder's remuneration.

Similarly, the standards by which we judge the "desirability" or "optimality" of an outcome will also depend on the setting. In evaluating public good choices, the criterion of "net social surplus" maximization is often invoked: does the public good decision maximize gross social benefit minus the cost of providing the goods? As for electing politicians, the property that a candidate would beat each competitor in head-to-head competition (i.e., would emerge a

---

[†] This article is a revised version of the lecture Eric S. Maskin delivered in Stockholm, Sweden, on December 8, 2007, when he received the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel. The article is copyright © The Nobel Foundation 2007 and is published here with the permission of the Nobel Foundation.

* School of Social Science, Institute for Advanced Study, Einstein Drive, Princeton, NJ 08540 (e-mail: maskin@ias.edu); Albert O. Hirschman Professor of Social Science, Institute for Advanced Study, and Visiting Lecturer, Princeton University. Research support from National Science Foundation grant SES–0318103 is gratefully acknowledged.

[1] There are many excellent surveys and textbook treatments of implementation theory that go into considerably more detail—both technical and conceptual—than I do here; see in particular: Andrew Postlewaite (1985), Theodore Groves and John Ledyard (1987), John Moore (1992), Thomas Palfrey (1992), chapter 10 of Martin Osborne and Ariel Rubinstein (1994), Beth Allen (1997), Luis Corchon (1996), Matthew Jackson (2001), Palfrey (2002), Roberto Serrano (2004), chapters 2 and 3 of David Austen-Smith and Jeffrey Banks (2005), chapter 6 of James Bergin (2005), chapters 14–16 of Allan Feldman and Serrano (2006), chapter 10 of Eric Rasmusen (2006), Sandeep Baliga and Tomas Sjöström (2007), and Corchon (2008). See also Partha Dasgupta, Peter Hammond, and Maskin (1979), Maskin and Sjöström (2002), Baliga and Maskin (2003), and my old survey, Maskin (1987).

Condorcet winner) is sometimes viewed as a natural desideratum (see Partha Dasgupta and Eric Maskin, forthcoming). In the auctioning of assets, there are two different criteria by which an outcome is typically judged: (a) whether the assets are put into the hands of bidders who value them the most (i.e., whether the allocation is *efficient*); and, alternatively, (b) whether the seller raises the greatest possible revenue from sales (i.e., whether *revenue maximization* is achieved). Finally, for the home buyer and builder, an outcome will ordinarily be considered "optimal" if it exhausts the potential gains from exchange between the parties, i.e., the house specification and remuneration are together Pareto optimal and individually rational.

A *mechanism* is an institution, procedure, or game for determining outcomes. Not surprisingly, who gets to choose the mechanism—i.e., who is the mechanism designer—will, once again, depend on the setting. In the case of public goods, we normally think of the government providing the goods as also choosing the method by which the levels of provision and financing are determined. Similarly, when it comes to sales of assets—where an *auction* is the typical mechanism—the asset seller often gets to call the shots about the rules, i.e., he is the one who chooses the auction format.

In the case of national political elections, by contrast, a mechanism is an *electoral procedure*, e.g., plurality rule, run-off voting, or the like. Moreover, the procedure is ordinarily prescribed long in advance, indeed sometimes by the country's constitution. Thus, here we should think of the framers of the constitution as the mechanism designers.

Finally, in the house-building example, a mechanism is a contract between the home buyer and builder and lays out the rights and responsibilities of each. Since these parties are presumably the ones who negotiate this contract, they themselves are the mechanism designers in this last setting.

Now, in the public framework, if the government knows at the outset which choice of public goods is optimal, then there is a simple—indeed, trivial—mechanism for achieving the optimum: the government has only to pass a law mandating this outcome. Similarly, if the auctioneer has prior knowledge of which bidders value the assets most, he can simply award them directly to those bidders (with or without payment).

The basic difficulty—which gives the subject of mechanism design its theoretical interest—is that the government or auctioneer will typically *not* have this information. After all, the net surplus-maximizing choice of public goods depends on citizens' *preferences* for such goods, and there is no particular reason why the government should know these preferences. Likewise, we wouldn't normally expect an auctioneer to know how much different potential buyers value the assets being sold.

Because mechanism designers do not generally know which outcomes are optimal in advance, they have to proceed more indirectly than simply prescribing outcomes by fiat; in particular, the mechanisms designed must generate the information needed as they are executed. The problem is exacerbated by the fact that the individuals who *do* have this critical information—the citizens in the public good case or the buyers in the asset-selling example—have their own objectives and so may not have the incentive to behave in a way that reveals what they know. Thus, the mechanisms must be *incentive compatible*. Much of the work in mechanism design, including my own, has been directed at answering three basic questions:

(A) When is it possible to design incentive-compatible mechanisms for attaining social goals?

(B) What form might these mechanisms take when they exist?

and

(C) When is finding such mechanisms ruled out theoretically?

That it is ever possible to design such mechanisms may, at first, seem surprising. How, after all, can a mechanism designer attain an optimal outcome without knowing exactly what he is aiming for? Thus, it may be helpful to consider a simple concrete example.

II. An Example

Consider a society consisting of two consumers of energy, Alice and Bob. An energy authority is charged with choosing the type of energy to be used by Alice and Bob. The options—from which the authority must make a single selection—are gas, oil, nuclear power, and coal.

Let us suppose that there are two possible states of the world. In state 1, the consumers place relatively little weight on the future, i.e., they have comparatively high temporal discount rates. In state 2, by contrast, they attach a great deal of importance to the future, meaning that their rates of discount are correspondingly low.

Alice, we will imagine, cares primarily about convenience when it comes to energy. This means that, in state 1, she will rank gas over oil, oil over coal, and coal over nuclear power, because as we move down her ranking, the energy source becomes either messier or more cumbersome to use. In state 2, by contrast, her ranking is

nuclear
gas
coal
oil

because she anticipates that technical advances will eventually make gas, coal, and especially nuclear power much easier to use—and, in this state, she lays particular stress on *future* benefits.

Bob is interested particularly in *safety*. This implies that in state 1, when he puts greatest weight on the present, he favors nuclear power over oil, oil over coal, and coal over gas. But if state 2 obtains—so that the future is comparatively important—his ranking is

oil
gas
coal
nuclear

which reflects the fact that, in the long run, the problem of disposing of nuclear waste can be expected to loom large, but that oil and gas safety are likely to improve somewhat.

To summarize, the consumers' rankings in the two states are given in Table 1.

Assume that the energy authority is interested in selecting an energy source that both consumers are reasonably happy with. If we interpret "reasonably happy" as getting one's first or second choice, then oil is the optimal choice in state 1, whereas gas is the best outcome in state 2. In the language of implementation theory, we say that the authority's *social choice rule* prescribes oil in state 1 and gas in state 2. Thus, if f is the social choice rule, it is given by Table 2.[2]

---

[2] In a more general setting, where Θ is the set of possible states of the world and A is the set of possible outcomes, a social choice rule f is a correspondence (a set-valued function) $f: \Theta \to\to A$, where, for any $\theta, f(\theta)$ is interpreted as the set of optimal outcomes in state θ (we are allowing for the possibility that more than one outcome might be considered optimal in a given state).

Table 1

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

Table 2

$f(\text{state } 1) = \text{oil}$	$f(\text{state } 2) = \text{gas}$

Suppose, however, that the authority does not know the state (although Alice and Bob do). This means that it does not know which alternative the social choice rule prescribes, i.e., whether oil or gas is the optimum.

Probably the most straightforward mechanism would be for the authority to ask each consumer to announce the state, whereupon it would choose oil if both consumers said "state 1," choose gas if both said "state 2," and flip a coin between them if it got a mixed response. But, notice that in this mechanism Alice has the incentive to say "state 2" regardless of the actual state and regardless of what Bob says, because she prefers gas to oil in both states. Indeed, by saying "state 2" rather than "state 1," she raises the probability of her preferred outcome from 0 to 0.5 if Bob says "state 1," and from 0.5 to 1 if Bob says "state 2." Hence, we would expect Alice to report "state 2" in both states. Similarly, Bob would always report "state 1," because he prefers oil to gas in either state. Taken together, Alice's and Bob's behavior implies that, in each state, the outcome is a 50–50 randomization between oil and gas. That is, there is only a 50 percent chance that the outcome is optimal, and so this mechanism is demonstrably too naïve.

Let us suppose, therefore, that the authority has the consumers participate in the mechanism given by Table 3:

Table 3

		Bob	
		Left	Right
Alice	Top	oil	coal
	Bottom	nuclear	gas

That is, Alice chooses "Top" or "Bottom" as her strategy; simultaneously, Bob chooses "Left" or "Right" as his strategy; and the outcome of those choices is given in the corresponding entry of the matrix.[3]

Observe that, in state 1, Bob is better off choosing Left regardless of what Alice does: if she plays Top, then Left leads to oil as the outcome (which Bob prefers), whereas Right gives rise to coal. If she plays Bottom, then nuclear power (Bob's preferred outcome) is the consequence of going Left, while Right leads to gas. That is, Left is the "dominant strategy" for Bob in state 1. Moreover, given that Bob is going Left, Alice is better off choosing Top rather than Bottom,

---

[3] More generally, a mechanism for a society with n individuals is a mapping $g: S_1 \times \cdots \times S_n \rightarrow A$, where, for all i, S_i is individual i's strategy space and $g(s_1, \ldots, s_n)$ is the outcome prescribed by the mechanism if individuals play the strategies $(s_1, \ldots, s_n)$.

because she prefers oil to nuclear power. Thus, in state 1, the clear prediction is for Alice to play Top and for Bob to play Left, i.e., (Top, Left) is the unique Nash equilibrium.[4] Furthermore—and this is the critical point—the resulting outcome, oil, is optimal in state 1.

Turning to state 2, we see that Bottom is the dominant strategy for Alice in that state. If Bob plays Left, then she is better off with Bottom than Top because she prefers nuclear power to oil. And if Bob goes Right, then Bottom leads to gas, which she prefers to the Top outcome, coal. With Alice choosing Bottom, Bob is better off going Right, because gas is better for him than nuclear power. Hence, in state 2, the (unique) Nash equilibrium is (Bottom, Right): Alice plays Bottom and Bob goes Right. Furthermore, this results in the optimal outcome, gas.

We have seen that in either state, the mechanism of Table 3 achieves the optimal outcome even though (a) the mechanism designer (the energy authority) does not even know the actual state, and (b) Alice and Bob are interested only in their own preferences, not those of the authority. More precisely, because the Nash equilibrium outcomes of the Table 3 mechanism coincide with the optimal outcomes in each state, we say that the mechanism *implements* the authority's social choice rule in Nash equilibrium.[5, 6]

III. A Brief History of Mechanism Design

The intellectual history of mechanism design theory goes back at least to nineteenth-century utopian socialists such as Robert Owen and Charles Fourier. Repulsed by what they viewed as the evils of the burgeoning capitalist system, these thinkers argued that socialism offered a more humane alternative and sometimes became involved in setting up experimental communities such as New Harmony, Indiana.

A more direct influence on the modern theory was the Planning Controversy, which reached its greatest intensity in the 1930s. The principal antagonists on one side were Oskar Lange and Abba Lerner, who argued forcefully that, done right, central planning could replicate the performance of free markets (Lange 1936 and Lerner 1944). Indeed, they suggested, planning could correct serious "market failures"—notably those on display in the Great Depression—and thereby potentially surpass markets. On the other side, Friedrich von Hayek and Ludwig von Mises staunchly denied the possibility that a planned system could ever approach the success of the free market (von Hayek 1944 and von Mises 1920).

The controversy was important and fascinating, but for certain onlookers such as Leonid Hurwicz, it was also rather frustrating. This was because it lacked conceptual precision: critical terms such as "decentralization" were left undefined. Moreover, the arguments adduced on either side often were often highly incomplete. In part, this was because they simply lacked the technical apparatus—in particular, game theory and mathematical programming—to generate truly persuasive conclusions.

This is where Leo Hurwicz entered the picture. Inspired by the debate, he attempted to provide unambiguous definitions of the central concepts, and this effort culminated in his two great

---

[4] In general, a Nash equilibrium is a specification of strategies—one for each individual—from which no individual has the incentive to deviate unilaterally. Thus, if $u_i(a,\theta)$ is individual i's payoff from outcome a in state θ, strategies $(s_1,\ldots,s_n)$ constitute a Nash equilibrium of mechanism g in state θ if $u_i(g(s_1,\ldots,s_i,\ldots,s_n),\theta) \geq u_i(g(s_1,\ldots,s_i',\ldots,s_n),\theta)$ for all i and all $s_i' \in S_i$.

[5] In a more general setting, mechanism g implements social choice rule f in Nash equilibrium if $f(\theta) = NE_g(\theta)$ for all θ, where $NE_g(\theta)$ is the set of Nash equilibrium outcomes of g in state θ.

[6] Nash equilibrium is a prediction of how individuals in a mechanism will behave. But a number of other predictive concepts—i.e., equilibrium concepts—have been considered in the implementation literature, among them subgame perfect equilibrium (Moore and Rafael Repullo 1988), undominated Nash equilibrium (Palfrey and Sanjay Srivastava 1991), Bayesian equilibrium (Postlewaite and David Schmeidler 1986), dominance solvability (Hervé Moulin 1979), trembling-hand perfect equilibrium (Sjöström 1993), and strong equilibrium (Bhaskar Dutta and Arunava Sen 1991).

papers, Hurwicz (1960) and (1972), where he also introduced the critical notion of incentive compatibility.

The work inspired by Hurwicz and others has produced a broad consensus among economists that von Hayek and von Mises were, in fact, correct—the market *is* the "best" mechanism—in settings where (a) there are large numbers of buyers and sellers, so that no single agent has significant market power; and (b) there are no significant externalities, that is, an agent's consumption, production, and information do not affect others' production or consumption.[7] However, mechanisms improving the market are generally possible if either assumption is violated.[8]

Hurwicz's work gave rise to an enormous literature, which has largely branched in two different directions. On the one hand, there is work that makes use of special, highly structured settings to study particular questions such as how to allocate public goods, how to design auctions, and how to structure contracts. On the other hand, there are studies obtaining results at a general, abstract level; that is, they make as few assumptions as possible about preferences, technologies, and so on. My own work has fallen into both categories at different times. But, in this paper, I will emphasize general results.

IV. Implementation of Social Choice Rules

Above I set out three central questions (A)–(C) about incentive-compatible mechanisms. Rephrased in the language of implementation theory, these questions become:

(A′) Under what conditions can a social choice rule be implemented?

(B′) What form does an implementing mechanism take?

(C′) Which social choice rules cannot be implemented?

In the mid-1970s I struggled with these questions. Eventually, I discovered that a property called *monotonicity* (now sometimes called Maskin-monotonicity) is the key to implementability in Nash equilibrium. Suppose that outcome a is optimal in state θ according to the social choice rule f in question, that is, $f(\theta) = a$. Then, if a doesn't fall in anyone's ranking relative to any other alternative in going from state θ to state θ', monotonicity requires that a also be optimal in state $\theta' : f(\theta') = a$. However, if *a does* fall relative to some outcome b in someone's ranking, monotonicity imposes no restriction.[9]

To see what monotonicity means more concretely, let's consider our energy example from before (see Tables 1 and 2). Recall that oil is the optimal outcome in state 1. Notice, too, that oil *falls* in Alice's ranking, relative to both coal and nuclear power, in going from state 1 to state 2 (Alice ranks oil higher than coal and nuclear in state 1, but just the opposite is true in state 2). Thus, the fact that gas—not oil—is optimal in state 2 does not violate monotonicity. Similarly, observe that gas falls in Bob's ranking, relative to both coal and nuclear power, in going from state 2 to state 1. Hence, even though gas is optimal in state 2, the fact that it is not optimal in

---

[7] See, for example, Peter Hammond (1979)—who shows, roughly, that the competitive market is the only incentive-compatible mechanism producing individually rational and Pareto optimal outcomes— and James Jordan (1982)—who shows the same thing when "incentive compatible" is replaced by "information efficient," under assumptions (i) and (ii).

[8] See, for instance, Theodore Groves (1973) and Edward Clarke (1971) for the case of public goods, and Jean-Jacques Laffont (1985) for the case of informational externalities.

[9] In a more general setting in which f can be set-valued, monotonicity requires that, for all states θ, θ' and all outcomes a, if $a \in f(\theta)$ and $u_i(a,\theta) \geq u_i(b,\theta)$ implies $u_i(a,\theta') \geq u_i(b,\theta')$ for all i and b, then $a \in f(\theta')$.

Table 4

State 1		State 2	
Alice	Bob	Alice	Bob
gas	nuclear	gas	nuclear
oil	oil	oil	oil
coal	coal	nuclear	coal
nuclear	gas	coal	gas
oil optimal		nuclear optimal	

state 1 is also not in conflict with monotonicity. Indeed, these verifications establish that the authority's social choice rule satisfies monotonicity (and thus the possibility of implementing it, which was shown earlier, does not contradict Theorem 1 below).

But suppose we modify the example somewhat, so that rankings and optimal outcomes are given by Table 4. With these changes, the social choice rule is no longer monotonic. Specifically, observe that although oil is optimal in state 1, it is not optimal in state 2, despite the fact that it falls in neither Alice's nor Bob's rankings between states 1 and 2 (given that oil doesn't fall, monotonicity would require it to remain optimal in state 2). Hence, we can conclude that there is *no* mechanism that implements the social choice rule of Table 4. More generally, we have:

THEOREM 1 (Maskin 1977): *If a social choice rule is implementable, then it must be monotonic.*

To see why the social choice rule in Table 4 is not implementable, suppose to the contrary that there *were* an implementing mechanism. Then, in particular, the mechanism would necessarily contain a pair of strategies (s_A, s_B)—for Alice and Bob, respectively—that result in outcome oil and constitute a Nash equilibrium in state 1.

I claim that (s_A, s_B) must also constitute a Nash equilibrium in state 2. To understand this claim, note first that Bob has no incentive to deviate unilaterally from s_B in state 2, since (i) he has no such incentive in state 1 (by definition of Nash equilibrium) and (ii) his preference ranking is the same in both states. Furthermore, Alice has no incentive to deviate from s_A in state 2. To see this, observe that if, contrary to the claim, Alice gained from deviating unilaterally from s_A in state 2, she must thereby be inducing the outcome gas (because this is the only outcome she prefers to oil in state 2). But Alice also prefers gas to oil in state 1, and so would benefit from the same deviation in that state, contradicting the assumption that (s_A, s_B) constitutes a Nash equilibrium in state 1.

Hence, (s_A, s_B) is indeed a Nash equilibrium in state 2. But the outcome it generates—oil—is not optimal in that state, establishing that the social choice rule is not implementable after all.

As we have seen, Tables 1 and 2 provide an example of a social choice rule that is monotonic and also implementable. However, it is not true that *all* monotonic social choice rules are implementable; see Maskin (1977) for a counterexample. Nevertheless, such counterexamples are rather contrived, and if an additional, often innocuous, condition is imposed, monotonicity *does* guarantee implementability, if there are at least three individuals in society.[10]

---

[10] That is not to say that implementation is impossible with just two individuals—indeed, our energy example of Tables 1 and 2 had only two individuals. However, as we will see below, implementation is facilitated by there being three or more individuals.

The additional condition is called *no veto power.* Suppose that all individuals, except possibly one, agree that a particular outcome *a* is *best*, meaning that they all put *a* at the *top* of their preference rankings. Then, if the social choice rule satisfies no veto power, *a* must be optimal. In other words, the remaining individual cannot "veto" it.

No veto power is especially innocuous—indeed, it imposes no restriction at all—when outcomes entail a distribution of economic goods across individuals. In that case, each individual will prefer a bigger share of those goods for himself or herself. So, no two of them can agree that a given outcome *a* is best: they cannot both get the biggest share. This means that, if there are three or more individuals, the hypothesis posited by the no veto power condition *cannot be satisfied*, and so logically the condition holds *automatically.*

A general result on the possibility of implementing social choice rules is the following:

THEOREM 2 (Maskin 1977): *Suppose that there are at least three individuals. If the social choice rule satisfies monotonicity and no veto power, then it is implementable.*

Proofs of Theorem 2 are beyond the scope of this paper (see Repullo 1987 for an especially elegant argument), but I should mention that they are usually *constructive.* That is, given the social choice rule to be implemented, a proof lays out an explicit recipe for the construction of a mechanism that does the trick.

It is worth pointing out why Theorem 2 posits at least three individuals. Often in economics, moving from two to three persons makes things more difficult.[11] But, for implementation theory, three individuals actually make matters easier. To understand why, remember that the underlying idea of a mechanism is to give individuals the incentive to behave in a way that ensures an optimal outcome. This entails "punishing" an individual for deviating from his prescribed (i.e., equilibrium) strategy. But if there are only two individuals, Alice and Bob, and one of them has deviated, it may be difficult to determine whether it was Alice who deviated and Bob who complied, or vice versa. This problem of identification is resolved once there are three people: a deviator sticks out more obviously when two or more other individuals are complying with equilibrium.

V. Concluding Remarks

This has been only a very brief introduction to implementation theory (which itself constitutes only part of the field of mechanism design). I have concentrated on work that was done over thirty years ago, which perhaps gives a misleadingly "antique" flavor to the paper. In fact, an especially gratifying aspect of the theory is that almost fifty years after Hurwicz (1960), the subject remains intellectually vibrant and important: new implementation papers are appearing all the time. It will be interesting to see where the field goes in the next fifty years.

REFERENCES

Allen, Beth. 1997. "Implementation Theory with Incomplete Information." In *Cooperation: Game Theoretic Approaches*, ed. S. Hart and A. Mas-Colell. Berlin: Springer.

Austen-Smith, David, and Jeffrey S. Banks. 2005. *Positive Political Theory II: Strategy and Structure.* Ann Arbor: University of Michigan Press.

---

[11] Zero-sum games provide a classic example of this phenomenon. The minimax theorem—which greatly simplifies the analysis of behavior in games—applies to *two*-person zero-sum games, but not, in general, to the case of *three* or more players.

Baliga, Sandeep, and Eric Maskin. 2003. "Mechanism Design for the Environment." In *Handbook of Environmental Economics,* Vol. 1, ed. K.-G. Mäler and J. R. Vincent, 305–24. Amsterdam: Elsevier Science, North-Holland.

Baliga, Sandeep, and Tomas Sjöström. 2007. "Mechanism Design: Recent Developments." In *The New Palgrave Dictionary of Economics*, ed. L. Blume and S. Durlauf. London: McMillan.

Bergin, James. 2005. *Microeconomic Theory: A Concise Course*. Oxford: Oxford University Press.

Clarke, Edward H. 1971. "Multipart Pricing of Public Goods." *Public Choice*, 11(1): 17–33.

Corchon, Luis C. 1996. *The Theory of Implementation of Socially Optimal Decisions in Economics*. London: Macmillan Press.

Corchon, Luis C. 2008. "The Theory of Implementation." In *The Encyclopedia of Complexity and System Science*. Berlin: Springer.

Dasgupta, Partha S., Peter J. Hammond, and Eric S. Maskin. 1979. "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility." *Review of Economic Studies*, 46(2): 185–216.

Dasgupta, Partha S., and Eric S. Maskin. Forthcoming. "On the Robustness of Majority Rule." *Journal of the European Economic Association*.

Dutta, Bhaskar, and Arunava Sen. 1991. "Implementation under Strong Equilibrium: A Complete Characterization." *Journal of Mathematical Economics*, 20(1): 46–67.

Feldman, Allan, and Roberto Serrano. 2006. *Welfare Economics and Social Choice Theory*. Berlin: Springer.

Groves, Theodore. 1973. "Incentives in Teams." *Econometrica*, 41(4): 617–31.

Groves, Theodore, and John O. Ledyard. 1987. "Incentive Compatibility since 1972." In *Information, Incentives, and Economic Mechanisms: Essays in Honor of Leonid Hurwicz*, ed. T. Groves, R. Radner, and S. Reiter, 48–111. Minneapolis: University of Minnesota Press.

Hammond, Peter J. 1979. "Straightforward Individual Incentive Compatibility in Large Economies." *Review of Economic Studies*, 46(2): 263–82.

Hurwicz, Leonid. 1960. "Optimality and Informational Efficiency in Resource Allocation Processes." In *Mathematical Methods in Social Sciences*, ed. Kenneth Arrow, S. Karlin and P. Suppes, 27–46. Stanford: Stanford University Press.

Hurwicz, Leonid. 1972. "On Informationally Decentralized Systems." In *Decision and Organization: A Volume in Honor of Jacob Marschak*, ed. C. B. McGuire and R. Radner, 297–336. Minneapolis: University of Minnesota Press.

Jackson, Matthew O. 2001. "A Crash Course in Implementation Theory." *Social Choice and Welfare*, 18(4): 655–708.

Jordan, James S. 1982. "The Competitive Allocation Process Is Informationally Efficient Uniquely." *Journal of Economic Theory*, 28(1): 1–18.

Laffont, Jean-Jacques M. 1985. "On the Welfare Analysis of Rational Expectations Equilibria with Asymmetric Information." *Econometrica*, 53(1): 1–29.

Lange, Oskar. 1936. "On the Economic Theory of Socialism." *Review of Economic Studies*, 4: 53–71.

Lerner, Abba. 1944. *The Economics of Control*. New York: McMillan.

Maskin, Eric S. 1977. "Nash Equilibrium and Welfare Optimality." *Review of Economic Studies*, 66(1): 23–38. (Published 1999.)

Maskin, Eric S. 1985. "The Theory of Implementation in Nash Equilibrium: A Survey." In *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, ed. L. Hurwicz, D. Schmeidler, and H. Sonnenschein, 173–204. Cambridge: Cambridge University Press.

Maskin, Eric S., and Tomas Sjöström. 2002. "Implementation Theory." In *Handbook of Social Choice and Welfare,* Vol. 1, ed. K. J. Arrow, A. K. Sen, and K. Suzumura, 237–88. Amsterdam: Elsevier Science, North-Holland.

Moore, John. 1992. "Implementation, Contracts, and Renegotiation in Environments with Complete Information." In *Advances in Economic Theory: Sixth World Congress Volume 1*, ed. Jean-Jacques Laffont, 182–282. Cambridge: Cambridge University Press.

Moore, John, and Rafael Repullo. 1988. "Subgame Perfect Implementation." *Econometrica*, 56(5): 1191–1220.

Moulin, Hervé. 1979. "Dominance Solvable Voting Schemes." *Econometrica*, 47(6): 1137–51.

Osborne, Martin J., and Ariel Rubinstein. 1994. *A Course in Game Theory.* Cambridge, MA: MIT Press.

Palfrey, Thomas R. 1992. "Implementation in Bayesian Equilibrium: The Multiple Equilibrium Problem in Mechanism Design." In *Advances in Economic Theory: Sixth World Congress Volume 1*, ed. Jean-Jacques Laffont, 283–323. Cambridge: Cambridge University Press.

Palfrey, Thomas R. 2001. "Implementation Theory." In *Handbook of Game Theory*, Vol. 3, ed. R. Aumann and S. Hart, 2271–2326. Amsterdam: Elsevier Science, North-Holland.

Palfrey, Thomas R., and Sanjay Srivastava. 1991. “Nash Implementation Using Undominated Strategies.” *Econometrica*, 59(2): 479–501.

Postlewaite, Andrew. 1985. “Implementation via Nash Equilibria in Economic Environments.” In *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, ed. L. Hurwicz, D. Schmeidler, and H. Sonnenschein, 205–28. Cambridge: Cambridge University Press.

Postlewaite, Andrew, and David Schmeidler. 1986. “Implementation in Differential Information Economies.” *Journal of Economic Theory*, 39(1): 14–33.

Rasmusen, Eric. 2006. *Games and Information: An Introduction to Game Theory*. Oxford: Blackwell Publishing.

Repullo, Rafael. 1987. “A Simple Proof of Maskin’s Theorem on Nash Implementation.” *Social Choice and Welfare*, 4(1): 39–41.

Serrano, Roberto. 2004. “The Theory of Implementation of Social Choice Rules.” *SIAM Review*, 46: 377–414.

Sjöström, Tomas. 1993. “Implementation in Perfect Equilibria.” *Social Choice and Welfare*, 10(1): 97–106.

von Hayek, Friedrich. 1944. *The Road to Serfdom*. London: Routledge.

von Mises, Ludwig. 1935. “Die Wirtschaftsrechnung im Sozialistischen Gemeinwesen.” In *Collectivist Economic Planning*, ed. F. von Hayek. London: Routledge.

This article has been cited by:

1. Nariaki Nishino, Koji Fukuya, Kanji Ueda. 2013. An Auction Mechanism Considering Seat Reservations in Movie Theater Services. *International Journal of Organizational and Collective Intelligence* **2**:1, 63-76. [CrossRef]
2. Susumu Cato. 2012. A note on the extension of a binary relation on a set to the power set. *Economics Letters* **116**:1, 46-48. [CrossRef]
3. Nariaki Nishino, Koji Fukuya, Kanji UedaAn Auction Mechanism Considering Seat Reservations in Movie Theater Services 193-211. [CrossRef]
4. Calvin Blackwell. 2011. Using a Simple Contest to Illustrate Mechanism Design. *The Journal of Economic Education* **42**:4, 375-387. [CrossRef]
5. Eric Smith, Martin Shubik. 2011. Endogenizing the provision of money: Costs of commodity and fiat monies in relation to the value of trade. *Journal of Mathematical Economics* . [CrossRef]
6. Barbara Krug, Hans Hendrischke. 2010. Market design in Chinese market places. *Asia Pacific Journal of Management* . [CrossRef]
7. John Merrifield. 2009. Defining Continuous Improvement and Cost Minimization Possibilities Through School Choice Experiments. *Journal of School Choice* **3**:3, 271-289. [CrossRef]
8. John Merrifield. 2009. Imagined Evidence and False Imperatives. *Journal of School Choice* **3**:1, 55-78. [CrossRef]
9. K. Ueda, T. Takenaka, J. Váncza, L. Monostori. 2009. Value creation and decision-making in sustainable society. *CIRP Annals - Manufacturing Technology* **58**:2, 681-700. [CrossRef]

The Evolution of Modern Portfolio Theory for the Institutional Investor

Bank of America Merrill Lynch

GLOBAL INSTITUTIONAL CONSULTING

SUMMER 2014

Ashvin B. Chhabra
Chief Investment Officer, Merrill Lynch Wealth Management
Head of Investment Management & Guidance

Most endowments and foundations use the policy portfolio as a guiding anchor for setting their investment strategy. The policy portfolio represents the synthesis of a strategic asset allocation anchored by the risk profile and return objectives of the institution, as well as a tactical overlay reflecting their best thinking of current market conditions and future returns.

Such an approach owes its origin to the work of Harry Markowitz[1] and others in the 1950s, now commonly referred to as Modern Portfolio Theory. Markowitz's work demonstrated the risk-adjusted benefits of portfolio diversification in an unambiguous manner. Today, the implementation of his technique of combining different sources of returns to reduce the overall risk of the portfolio is simply known as asset allocation.

The early pioneers of Endowment Investing[2], notably David Swensen and Jack Meyer, correctly foresaw that certain investment structures and securities could provide different risk-return characteristics than those provided by standard asset classes. The addition of absolute return structures, illiquid private equity, venture capital and timber was a logical evolution of Markowitz's insight on portfolio diversification. The policy portfolio can thus be viewed as asset allocation on steroids, albeit with additional insight and non-trivial implementation hurdles (see Exhibit 1). The early adopters of Endowment Investing were also able to successfully deliver excess return (alpha) through astute manager and security selection. However, there is a clear tradeoff between maximizing return and diversifying the overall portfolio.

> "The policy portfolio can thus be viewed as asset allocation on steroids, albeit with additional insight and non-trivial implementation hurdles."

Exhibit 1: The Policy Portfolio is Asset Allocation on Steroids

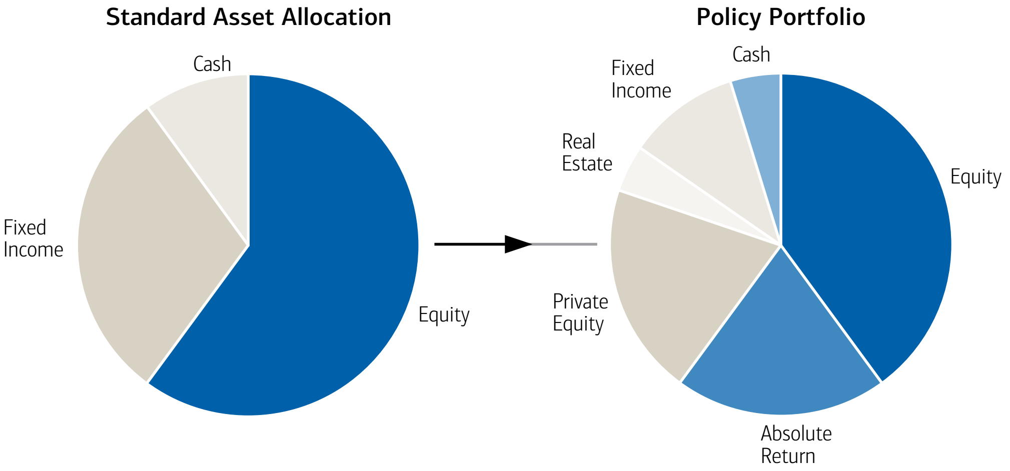

Source: GWM Investment Management & Guidance

Concentrating one's chips on a few managers that have the promise to maximize alpha may be the theoretically pure choice, but not without significant downside. Boston University's ill-fated investment in Seragen in the 1980s serves as an important reminder of the dangers of over-confidence and concentration. Therefore, most endowments diversify prudently, thus inevitably reducing both idiosyncratic risk and expected excess return.

The Problems of Trend Chasing and Herd Behavior

Nevertheless, the notion of delivering superior or even out-sized returns for one's institution through a differentiated investment strategy is a powerful one. In a recent paper, Goetzmann[3] demonstrates the hypnotic effect this philosophy has had on Endowment Investing – leading to trend-chasing and herd behavior. In particular, he found that endowments with below-median allocations to alternative investments increased their allocations to hedge funds in an effort to catch up with their "school's nearest competitors". He argues that "the dynamic patterns that we see in asset allocations in university endowments are consistent with an arms race model of universities."

> “ the dynamic patterns that we see in asset allocations in university endowments are consistent with an arms race model of universities...”

A fair question to ask is whether this kind of trend-chasing and herd behavior led to the exacerbation of the problems experienced by endowments and foundations in 2008.

The purpose of this article is to stress that the source of those problems is much deeper than one might imagine and has its roots in how we understand Modern Portfolio Theory in the context of institutional investing. Given the continued economic uncertainty and fragility of financial markets, I hope that the formalism discussed below can be helpful in setting an appropriate investment strategy for institutions.

The Greater Objective -- Meeting Goals

To begin, one must note that maximizing returns through any particular investment strategy is merely part of a greater objective of meeting a series of (often conflicting!) goals for the institution. *The most immediate and essential of these goals is to provide cash flows, as much and when needed, for an individual or institution to function effectively.* Thus both the Markowitz framework (asset allocation) and the Endowment Model (policy portfolio) need to be understood within a larger framework of institutional needs.

An Important Step: Bond Immunization

An important step in connecting the diversified market portfolio with the particular needs of the investing institution was taken by James Tobin[4] a few years after Markowitz wrote his famous paper. Tobin pointed out that corporations had well-defined liabilities that needed to be met with certainty. This certainty was at odds with the uncertainties of even a well-diversified market portfolio. Tobin's solution involved creating a bond portfolio that matched the near-term liabilities. This technique is now known as bond immunization. Through duration matching, the coupons of the bond portfolio match the firm's short to medium term liabilities and protect those liabilities from market fluctuations.

Two Distinct Approaches: Risky and Riskless Portfolios

In the broader context this step led to a framework that combined two distinct approaches—asset-liability management (the riskless portfolio) and the standard diversified market (risky) portfolio. Of course, the two distinct portfolios can be visually represented by a single pie-chart, but the distinction is of critical importance.

Why is the distinction so important?

Detailed historical studies of global financial markets[5,6] show quite clearly that they are susceptible to instabilities of all magnitudes. Markets, as Mandelbrot[7] emphasized, are not just mildly random but wildly random.

In return distributions these instabilities are evident as fat tails. However, the fat tails do not do justice to these instabilities as such events do not occur randomly. Rather, these instabilities are highly correlated and come in clusters. In extreme cases, some financial markets simply go out of existence for extended periods of time.

The riskless portfolio concept thus deserves serious consideration. The generalization of bond immunization is immunization against all the kinds of risks that an institution can face (whenever possible or affordable).

The Challenge of Idiosyncratic Risk and Return

A second point worthy of further discussion is the impact of idiosyncratic risk and return. CIOs spend a great deal of effort selecting managers that provide the promise of significant alpha to the portfolio while at the same time judiciously combining them to create a diversified portfolio. However, the role (and impact) of alpha in the Endowment Model has been fundamentally misunderstood by the slew of new adopters.

Let us consider what impact alpha, within the context of a diversified portfolio, can have on an institution.

First, let us examine the increase in alpha that can be added by superior performance, relative to an already sophisticated peer group (see Exhibit 2).

Given this data[8], we can ask the following question: What impact can this superior performance have on the size of the total endowment over time?

To answer this question, it is instructive to examine the size distribution of the list of endowments covered by NACUBO in its annual survey (see Exhibit 3).

Let us now assume that an institution of median size (50th percentile) would like to grow its endowment aggressively. Assume further that, by investing aggressively in illiquid and alternative strategies (within the diversified framework), it is able to systematically deliver top quartile performance. At the same time consider an endowment in the 75th percentile in terms of size that is sluggish in its strategy and only manages bottom quartile performance. Let them do this year after year.

Neglecting any additional risks arising from this out-performing strategy adopted by the smaller endowment, let us compute how many years it takes before the median sized endowment breaks into the top quartile in terms of size.

Exhibit 2:
NACUBO-Common Fund Study of Endowment Performance (2013)

Endowment Performance (%)	10 Year Return (Annualized)
25%	6.3%
50%	7.1%
75%	7.8%

Source: NACUBO – Common Fund Study of Endowment Performance (2013)

Exhibit 3:
Distribution of Endowment Sizes

Endowment Size (%)	Endowment Size
25%	$42MM
50%	$101MM
75%	$337MM
90%	$965MM
95%	$1,733MM

Source: NACUBO – Common Fund Study of Endowment Performance (2013)

The smaller institution will catch up (even under these very favorable assumptions) in about 65 years. The time frame is similar for an institution in the 75th percentile chasing one already in the 90th percentile.

In retrospect the answer is obvious. The distribution of endowment sizes is not normal but instead has a very long tail. Consistent out-performance is hard to come by and it is hard to move the needle unless one takes a lot of idiosyncratic risk and veers off the beaten path. This observation is consistent with what is seen when studying endowments, foundations or the net-worth of individual investors[9].

Endowments may have a long-term view of their institution, but the day-to-day reality of the investment process has a much shorter time frame. Studies have shown that after three years of under-performance CIOs exit a manager or strategy. In summary, an aggressive investment strategy stuck within a diversified framework is not the ticket to wealth mobility, or to breaking out of one's peer group. This is true for both individual investors and for institutions. However, the desire to provide top quartile performance within an already sophisticated peer group can lead to some pretty skewed strategies – leading many to disregard the Hippocratic admonition of first doing no harm.

Bucket 1:

A Safety Portfolio: A risk mitigation portfolio that may include Asset Liability Management and other protective strategies including tail-risk hedging.

Bucket 2:

A Policy Portfolio: A diversified market portfolio where the source of return is driven by various market betas with an overlay of some alpha. This includes hedge funds and private equity.

Bucket 3:

An Aspirational Portfolio: This portfolio consists of a collection of investments that may have a disproportionately positive impact on the institution, usually accompanied by much higher risk. These investments may be concentrated around areas where the institution has or can build a sustainable advantage such as building a top-notch medical facility, sports team or fund-raising operation.

A Third Portfolio Bucket

In order to better analyze the overall portfolio, it turns out to be important and useful to separate investments that have exceedingly high return expectations into a third portfolio bucket, as they come with a risk-return profile different from that of a diversified market portfolio.[10]

The three portfolios together create a framework for understanding the entire portfolio in terms of objectives, risk and return within the context of fat tails and black swans (see Exhibit 4).

Exhibit 4: The Wealth Allocation Framework for Institutional Investors

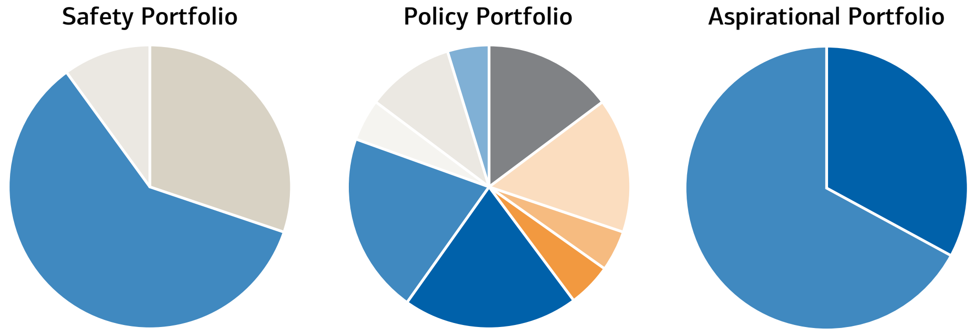

Source: GWM Investment Management & Guidance

The Critical Role of Risk Allocation

The critical piece in this three-portfolio framework is the concept of Risk Allocation.

Risk Allocation is a fundamentally more important concept than asset allocation[11]. In fact asset allocation is simply a special case of Risk Allocation. The policy portfolio is a piece, albeit an important one, within the overall framework. The optimal Risk Allocation involves balancing allocations among the three (inter-related) risk buckets and must be set in the context of meeting the goals and objectives of the institution.

Failure to understand Risk Allocation can lead to structural imbalances in the investment portfolio, as evidenced by the liquidity crisis in endowment portfolios during the 2008 crash. The following example may be illustrative:

In the quest to increase returns, many endowments embarked on an aggressive program to capture the illiquidity premium through investments in private equity. In 2008, the policy portfolio of several institutions indicated that illiquid securities ranged from a third to a half of their holdings. Furthermore, these institutions had PE commitments approximating about 10% a year over the next three years. These commitments were either to be balanced by expected distributions or were part of an overall strategy to increase exposure to alternative investments. While this situation may look like a sensible diversified investment strategy in terms of the policy portfolio, it is instructive to look at it in the context of the Wealth Allocation Framework (see Exhibit 5).

Exhibit 5: Evolution of the Policy portfolio?

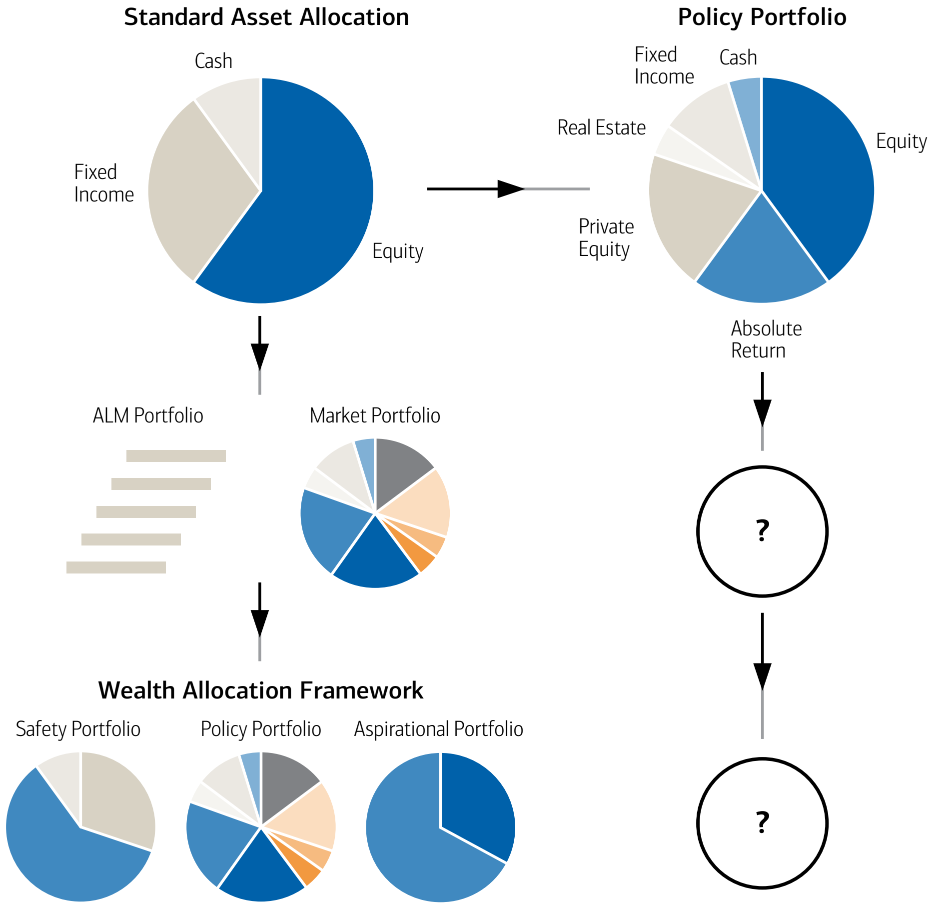

Source: GWM Investment Management & Guidance

The Wealth Allocation Framework for Institutional Investors

The Wealth Allocation Framework forces the portfolio to separate out the risk-mitigation assets from market assets and aspirational assets. One then adds the liabilities associated with contributions to the operating budget and commitments to private equity capital calls.

Since private equity commitments are firm commitments, they force one to debit the safety portfolio/bucket (rather than the market portfolio) by the same amount. This is not true for other non-recourse leverage that may be tied to a particular investment. Hence the critical distinction between recourse and non-recourse leverage is evident in this framework.

Most institutions were highly under-allocated to bonds and cash (Harvard's 2008 cash allocation was -5%). This led to an overall net negative safety portfolio for many institutions. The equity and equity-like market portfolio had an approximately 50% allocation to illiquid assets. When the total value of the market portfolio decreased by 30% to 40% (a major instability or market crash) illiquid investments as a percentage of the total balance sheet went up sharply. This was accentuated by the fact that the dollar amount of the PE commitments did not go down at all. This de-correlation between the size of the PE commitments and that of the PE portfolio itself is evident in the Wealth Allocation Framework (see Exhibit 6).

Exhibit 6: A Risk Balance Sheet Approach

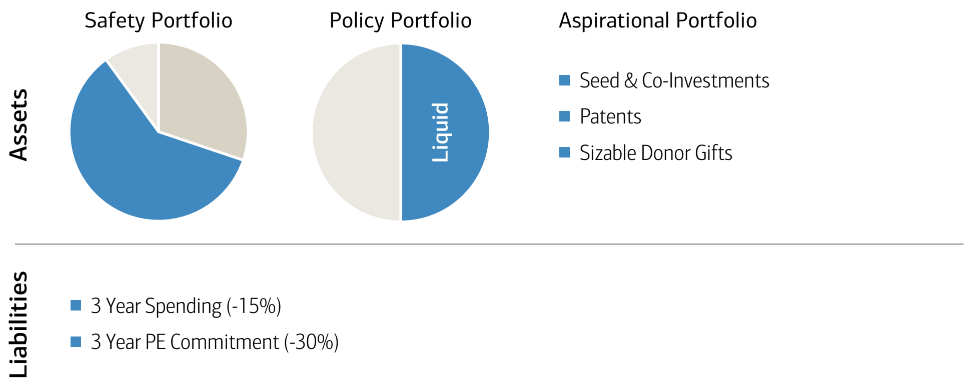

Source: GWM Investment Management & Guidance

The emergent picture clearly indicates that such a risk allocation cannot withstand a major market downturn without severe stress.

Additional Comments

Exposure to Beta

There is a choice of how much exposure to equity markets one may want in the market portfolio. The total exposure can also be quantified as a single number, beta (sensitivity to the equity market), in order to incorporate other sources of return (such as credit spreads) that may be correlated to equity markets. It is somewhat surprising to see that many institutions with widely varying risk profiles and dependencies on their endowments all seem to have similar betas ranging from 0.6 to 0.75. This situation may have its origin in the common historical starting point for most endowments and foundations, the conventional 60/40 stock and bond portfolio. For an understanding of what happens to these high-beta portfolios that contain a large percentage of alternative assets in extreme market conditions, the concept of stress beta introduced by Leibowitz,[12] is important and especially useful.

Risk Parity

Risk Parity is a special case of risk allocation. However, one may ask, "parity with respect to what?" In some cases it may be useful to think about a target volatility to compare different sources of return while building a portfolio. In the broader context, however, it may be more useful for an institution to develop the right risk sensitivities or allocation with regard to a series of possible economic and market scenarios that the portfolio or institution may be exposed to.

In the end the institution's investment strategy is simply part of a much broader effort to control the risk of not achieving institutional goals under a variety of possible market scenarios. The Wealth Allocation Framework clarifies the role of an investment strategy within the broader context of essential and aspirational goals and stable and unstable economic scenarios.

The Role of True Alpha

True alpha comes from idiosyncratic investing—tactical asset allocation, security selection and so on. However, for it to make a difference in the aggregate to the institution, this strategy must be accomplished on a large scale i.e. on a scale comparable to the institution's wealth. This is neither desirable nor feasible for most endowments and is accomplished only occasionally. Successful examples are few and far between and in most cases, it means charting one's own course and having the backing of the institution to stay the course for many decades[13].

As a rule, idiosyncratic return from superior investment and manager selection will only contribute a small percentage to the total wealth of an institution. Therefore, the risks taken to achieve these excess returns should be small in the aggregate. Not every investment portfolio can execute well on the Aspirational Portfolio. Instead, an institution may try to fill this bucket through other activities such as raising large donor gifts or identifying intellectual property that may yield potentially lucrative patents. Other examples may be developing a top-notch medical facility or a top-tier sports team. The overall portfolio construction should focus on achieving institutional goals. In the institutional sense, this demands a closer connection between the investment strategy and other activities of the institution.

Acknowledgements

I would like to thank Nevenka Vrdoljak and Ravindra Koneru for assistance with the figures and calculations for this paper. An earlier version of this paper was published in the NMS Exchange newsletter in November 2012.

Sources, References and Further Reading

[1] Harry Markowitz, The Journal of Finance, Vol. 7, No. 1. (1952)

[2] David F Swensen: Pioneering Portfolio Management, Published by Free Press (2000)

[3] William Goetzmann & Sharon Oster, Competition Among University Endowments, Yale ICF Working Paper No. 06-14, (2012)

[4] J. Tobin, Liquidity preference as behavior toward risk. Review of Economic Studies 25 (1958)

[5] Carmen M. Reinhart & Kenneth S. Rogoff, This Time Is Different: Eight Centuries of Financial Folly, Published by Princeton University Press (2011)

[6] Elroy Dimson, Paul Marsh, and Mike Staunton, Triumph of the Optimists: 101 Years of Global Investment Returns. Published by Princeton University Press (2002)

[7] B.B. Mandelbrot, The Random Character of Stock Market Prices ed. P Cootner, Published by MIT Press (1964)

[8] Nacubo-Commonfund Study 2013 (see *www.nacubo.org*)

[9] Ashvin B. Chhabra, Ravindra Koneru & Lex Zaharoff, “Modern Portfolio Theory’s Third Rail”, Journal of Wealth Management, Vol. 14, No. 1 (2011)

[10] Ashvin B. Chhabra, “Beyond Markowitz: A comprehensive Wealth Allocation Framework for individual investors”, Journal of Wealth Management, Vol. 7, No. 4 (2005)

[11] Ashvin B. Chhabra, “Clarifications on Beyond Markowitz”, Journal of Wealth Management, Vol. 11, No. 3 (2007)

[12] Martin L. Leibowitz, Anthony Bova, & P. Brett Hammond, The Endowment Model of Investing: Return, Risk, and Diversification, Published by Wiley Finance (2010)

[13] David Chambers and Elroy Dimson, Keynes the stock market investor, available at *http://ssrn.com/abstract=2023011*

GWM Investment Management & Guidance (IMG) provides industry-leading investment solutions, portfolio construction advice and wealth management guidance. This material was prepared by the Investment Management & Guidance Group (IMG) and is not a publication of BofA Merrill Lynch Global Research. The views expressed are those of IMG only and are subject to change. This information should not be construed as investment advice. It is presented for information purposes only and is not intended to be either a specific offer by any Merrill Lynch entity to sell or provide, or a specific invitation for a consumer to apply for, any particular retail financial product or service that may be available.

This information and any discussion should not be construed as a personalized and individual client recommendation, which should be based on each client's investment objectives, risk tolerance, liquidity needs and financial situation. This information and any discussion also is not intended as a specific offer by Merrill Lynch, its affiliates, or any related entity to sell or provide, or a specific invitation for a consumer to apply for, any particular retail financial product or service. Investments and opinions are subject to change due to market conditions and the opinions and guidance may not be profitable or realized. Any information presented in connection with BofA Merrill Lynch Global Research is general in nature and is not intended to provide personal investment advice. The information does not take into account the specific investment objectives, financial situation and particular needs of any specific person who may receive it. Investors should understand that statements regarding future prospects may not be realized.

Asset allocation and diversification do not assure a profit or protect against a loss during declining markets.

© 2014 Bank of America Corporation. All rights reserved.

Bank of America
Merrill Lynch

The Evolution of Modern Portfolio Theory for the Institutional Investor

Bank of America Merrill Lynch

GLOBAL INSTITUTIONAL CONSULTING

SUMMER 2014

Ashvin B. Chhabra
Chief Investment Officer, Merrill Lynch Wealth Management
Head of Investment Management & Guidance

Most endowments and foundations use the policy portfolio as a guiding anchor for setting their investment strategy. The policy portfolio represents the synthesis of a strategic asset allocation anchored by the risk profile and return objectives of the institution, as well as a tactical overlay reflecting their best thinking of current market conditions and future returns.

Such an approach owes its origin to the work of Harry Markowitz[1] and others in the 1950s, now commonly referred to as Modern Portfolio Theory. Markowitz's work demonstrated the risk-adjusted benefits of portfolio diversification in an unambiguous manner. Today, the implementation of his technique of combining different sources of returns to reduce the overall risk of the portfolio is simply known as asset allocation.

The early pioneers of Endowment Investing[2], notably David Swensen and Jack Meyer, correctly foresaw that certain investment structures and securities could provide different risk-return characteristics than those provided by standard asset classes. The addition of absolute return structures, illiquid private equity, venture capital and timber was a logical evolution of Markowitz's insight on portfolio diversification. The policy portfolio can thus be viewed as asset allocation on steroids, albeit with additional insight and non-trivial implementation hurdles (see Exhibit 1). The early adopters of Endowment Investing were also able to successfully deliver excess return (alpha) through astute manager and security selection. However, there is a clear tradeoff between maximizing return and diversifying the overall portfolio.

> “The policy portfolio can thus be viewed as asset allocation on steroids, albeit with additional insight and non-trivial implementation hurdles.”

Exhibit 1: The Policy Portfolio is Asset Allocation on Steroids

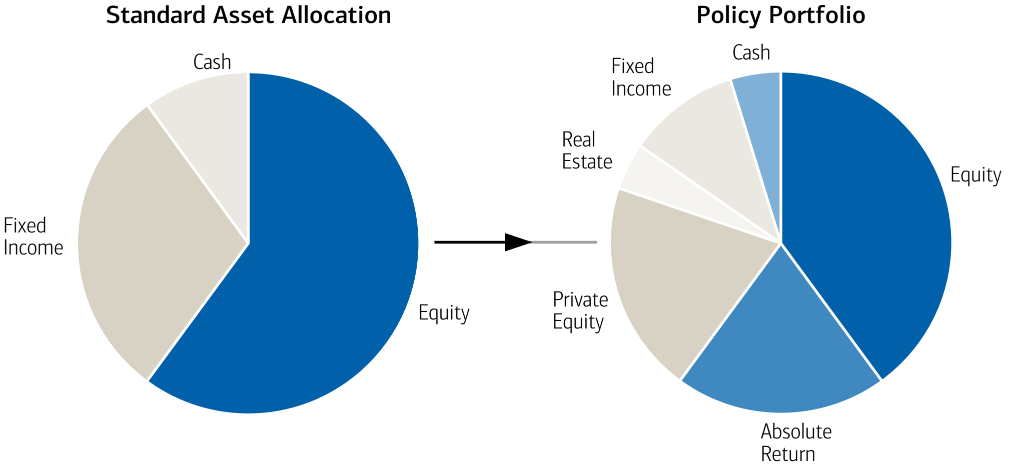

Source: GWM Investment Management & Guidance

Concentrating one's chips on a few managers that have the promise to maximize alpha may be the theoretically pure choice, but not without significant downside. Boston University's ill-fated investment in Seragen in the 1980s serves as an important reminder of the dangers of over-confidence and concentration. Therefore, most endowments diversify prudently, thus inevitably reducing both idiosyncratic risk and expected excess return.

The Problems of Trend Chasing and Herd Behavior

Nevertheless, the notion of delivering superior or even out-sized returns for one's institution through a differentiated investment strategy is a powerful one. In a recent paper, Goetzmann[3] demonstrates the hypnotic effect this philosophy has had on Endowment Investing – leading to trend-chasing and herd behavior. In particular, he found that endowments with below-median allocations to alternative investments increased their allocations to hedge funds in an effort to catch up with their "school's nearest competitors". He argues that "the dynamic patterns that we see in asset allocations in university endowments are consistent with an arms race model of universities."

> “ the dynamic patterns that we see in asset allocations in university endowments are consistent with an arms race model of universities...”

A fair question to ask is whether this kind of trend-chasing and herd behavior led to the exacerbation of the problems experienced by endowments and foundations in 2008.

The purpose of this article is to stress that the source of those problems is much deeper than one might imagine and has its roots in how we understand Modern Portfolio Theory in the context of institutional investing. Given the continued economic uncertainty and fragility of financial markets, I hope that the formalism discussed below can be helpful in setting an appropriate investment strategy for institutions.

The Greater Objective -- Meeting Goals

To begin, one must note that maximizing returns through any particular investment strategy is merely part of a greater objective of meeting a series of (often conflicting!) goals for the institution. *The most immediate and essential of these goals is to provide cash flows, as much and when needed, for an individual or institution to function effectively.* Thus both the Markowitz framework (asset allocation) and the Endowment Model (policy portfolio) need to be understood within a larger framework of institutional needs.

An Important Step: Bond Immunization

An important step in connecting the diversified market portfolio with the particular needs of the investing institution was taken by James Tobin[4] a few years after Markowitz wrote his famous paper. Tobin pointed out that corporations had well-defined liabilities that needed to be met with certainty. This certainty was at odds with the uncertainties of even a well-diversified market portfolio. Tobin's solution involved creating a bond portfolio that matched the near-term liabilities. This technique is now known as bond immunization. Through duration matching, the coupons of the bond portfolio match the firm's short to medium term liabilities and protect those liabilities from market fluctuations.

Two Distinct Approaches: Risky and Riskless Portfolios

In the broader context this step led to a framework that combined two distinct approaches—asset-liability management (the riskless portfolio) and the standard diversified market (risky) portfolio. Of course, the two distinct portfolios can be visually represented by a single pie-chart, but the distinction is of critical importance.

Why is the distinction so important?

Detailed historical studies of global financial markets[5,6] show quite clearly that they are susceptible to instabilities of all magnitudes. Markets, as Mandelbrot[7] emphasized, are not just mildly random but wildly random.

In return distributions these instabilities are evident as fat tails. However, the fat tails do not do justice to these instabilities as such events do not occur randomly. Rather, these instabilities are highly correlated and come in clusters. In extreme cases, some financial markets simply go out of existence for extended periods of time.

The riskless portfolio concept thus deserves serious consideration. The generalization of bond immunization is immunization against all the kinds of risks that an institution can face (whenever possible or affordable).

The Challenge of Idiosyncratic Risk and Return

A second point worthy of further discussion is the impact of idiosyncratic risk and return. CIOs spend a great deal of effort selecting managers that provide the promise of significant alpha to the portfolio while at the same time judiciously combining them to create a diversified portfolio. However, the role (and impact) of alpha in the Endowment Model has been fundamentally misunderstood by the slew of new adopters.

Let us consider what impact alpha, within the context of a diversified portfolio, can have on an institution.

First, let us examine the increase in alpha that can be added by superior performance, relative to an already sophisticated peer group (see Exhibit 2).

Given this data[8], we can ask the following question: What impact can this superior performance have on the size of the total endowment over time?

To answer this question, it is instructive to examine the size distribution of the list of endowments covered by NACUBO in its annual survey (see Exhibit 3).

Let us now assume that an institution of median size (50th percentile) would like to grow its endowment aggressively. Assume further that, by investing aggressively in illiquid and alternative strategies (within the diversified framework), it is able to systematically deliver top quartile performance. At the same time consider an endowment in the 75th percentile in terms of size that is sluggish in its strategy and only manages bottom quartile performance. Let them do this year after year.

Neglecting any additional risks arising from this out-performing strategy adopted by the smaller endowment, let us compute how many years it takes before the median sized endowment breaks into the top quartile in terms of size.

Exhibit 2:
NACUBO-Common Fund Study of Endowment Performance (2013)

Endowment Performance (%)	10 Year Return (Annualized)
25%	6.3%
50%	7.1%
75%	7.8%

Source: NACUBO – Common Fund Study of Endowment Performance (2013)

Exhibit 3:
Distribution of Endowment Sizes

Endowment Size (%)	Endowment Size
25%	$42MM
50%	$101MM
75%	$337MM
90%	$965MM
95%	$1,733MM

Source: NACUBO – Common Fund Study of Endowment Performance (2013)

The smaller institution will catch up (even under these very favorable assumptions) in about 65 years. The time frame is similar for an institution in the 75th percentile chasing one already in the 90th percentile.

In retrospect the answer is obvious. The distribution of endowment sizes is not normal but instead has a very long tail. Consistent out-performance is hard to come by and it is hard to move the needle unless one takes a lot of idiosyncratic risk and veers off the beaten path. This observation is consistent with what is seen when studying endowments, foundations or the net-worth of individual investors[9].

Endowments may have a long-term view of their institution, but the day-to-day reality of the investment process has a much shorter time frame. Studies have shown that after three years of under-performance CIOs exit a manager or strategy. In summary, an aggressive investment strategy stuck within a diversified framework is not the ticket to wealth mobility, or to breaking out of one's peer group. This is true for both individual investors and for institutions. However, the desire to provide top quartile performance within an already sophisticated peer group can lead to some pretty skewed strategies – leading many to disregard the Hippocratic admonition of first doing no harm.

Bucket 1:

A Safety Portfolio: A risk mitigation portfolio that may include Asset Liability Management and other protective strategies including tail-risk hedging.

Bucket 2:

A Policy Portfolio: A diversified market portfolio where the source of return is driven by various market betas with an overlay of some alpha. This includes hedge funds and private equity.

Bucket 3:

An Aspirational Portfolio: This portfolio consists of a collection of investments that may have a disproportionately positive impact on the institution, usually accompanied by much higher risk. These investments may be concentrated around areas where the institution has or can build a sustainable advantage such as building a top-notch medical facility, sports team or fund-raising operation.

A Third Portfolio Bucket

In order to better analyze the overall portfolio, it turns out to be important and useful to separate investments that have exceedingly high return expectations into a third portfolio bucket, as they come with a risk-return profile different from that of a diversified market portfolio.[10]

The three portfolios together create a framework for understanding the entire portfolio in terms of objectives, risk and return within the context of fat tails and black swans (see Exhibit 4).

Exhibit 4: The Wealth Allocation Framework for Institutional Investors

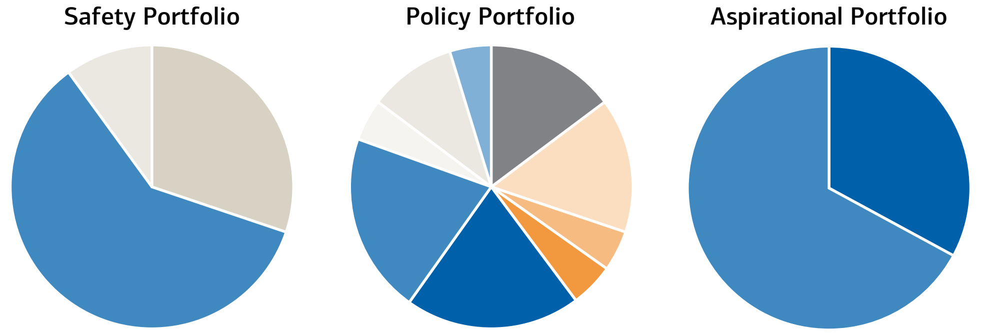

Source: GWM Investment Management & Guidance

The Critical Role of Risk Allocation

The critical piece in this three-portfolio framework is the concept of Risk Allocation.

Risk Allocation is a fundamentally more important concept than asset allocation[11]. In fact asset allocation is simply a special case of Risk Allocation. The policy portfolio is a piece, albeit an important one, within the overall framework. The optimal Risk Allocation involves balancing allocations among the three (inter-related) risk buckets and must be set in the context of meeting the goals and objectives of the institution.

Failure to understand Risk Allocation can lead to structural imbalances in the investment portfolio, as evidenced by the liquidity crisis in endowment portfolios during the 2008 crash. The following example may be illustrative:

In the quest to increase returns, many endowments embarked on an aggressive program to capture the illiquidity premium through investments in private equity. In 2008, the policy portfolio of several institutions indicated that illiquid securities ranged from a third to a half of their holdings. Furthermore, these institutions had PE commitments approximating about 10% a year over the next three years. These commitments were either to be balanced by expected distributions or were part of an overall strategy to increase exposure to alternative investments. While this situation may look like a sensible diversified investment strategy in terms of the policy portfolio, it is instructive to look at it in the context of the Wealth Allocation Framework (see Exhibit 5).

Exhibit 5: Evolution of the Policy portfolio?

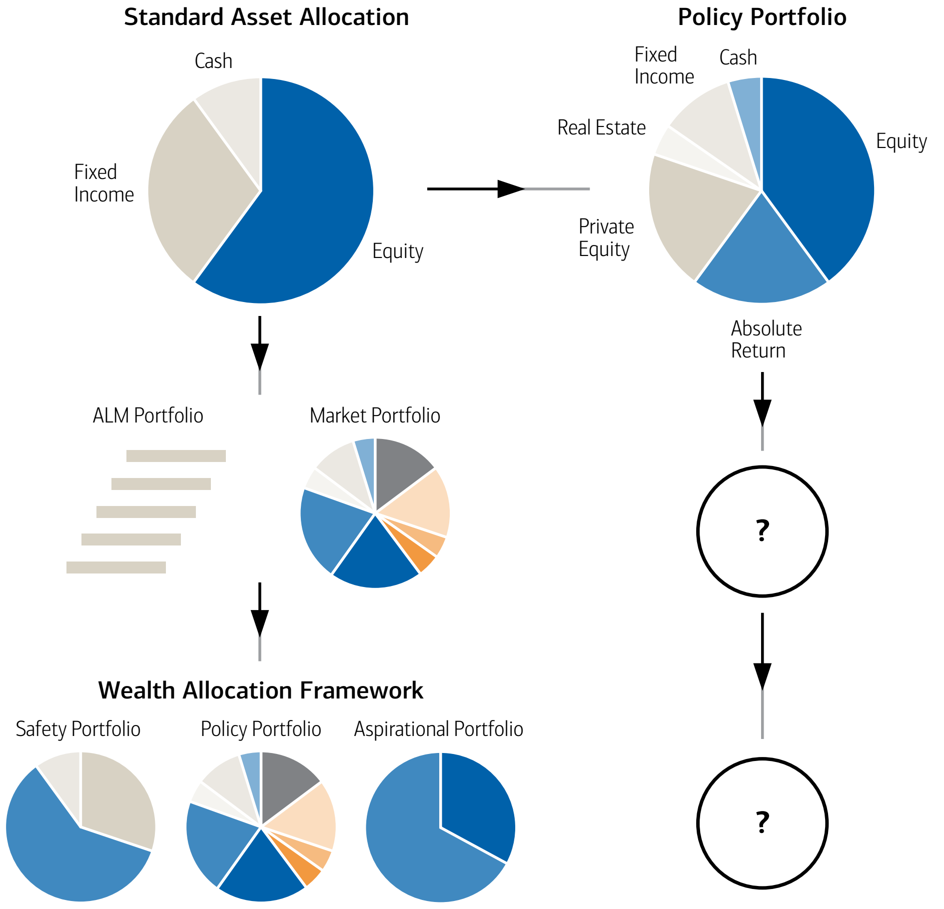

Source: GWM Investment Management & Guidance

The Wealth Allocation Framework for Institutional Investors

The Wealth Allocation Framework forces the portfolio to separate out the risk-mitigation assets from market assets and aspirational assets. One then adds the liabilities associated with contributions to the operating budget and commitments to private equity capital calls.

Since private equity commitments are firm commitments, they force one to debit the safety portfolio/bucket (rather than the market portfolio) by the same amount. This is not true for other non-recourse leverage that may be tied to a particular investment. Hence the critical distinction between recourse and non-recourse leverage is evident in this framework.

Most institutions were highly under-allocated to bonds and cash (Harvard's 2008 cash allocation was -5%). This led to an overall net negative safety portfolio for many institutions. The equity and equity-like market portfolio had an approximately 50% allocation to illiquid assets. When the total value of the market portfolio decreased by 30% to 40% (a major instability or market crash) illiquid investments as a percentage of the total balance sheet went up sharply. This was accentuated by the fact that the dollar amount of the PE commitments did not go down at all. This de-correlation between the size of the PE commitments and that of the PE portfolio itself is evident in the Wealth Allocation Framework (see Exhibit 6).

Exhibit 6: A Risk Balance Sheet Approach

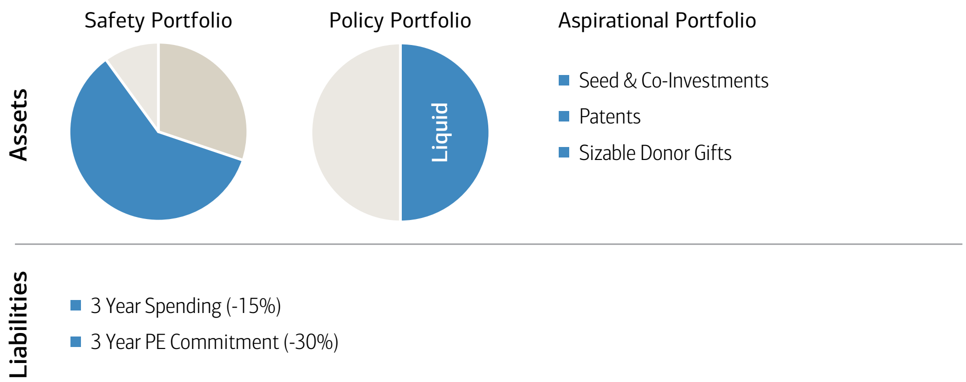

Source: GWM Investment Management & Guidance

The emergent picture clearly indicates that such a risk allocation cannot withstand a major market downturn without severe stress.

Additional Comments

Exposure to Beta

There is a choice of how much exposure to equity markets one may want in the market portfolio. The total exposure can also be quantified as a single number, beta (sensitivity to the equity market), in order to incorporate other sources of return (such as credit spreads) that may be correlated to equity markets. It is somewhat surprising to see that many institutions with widely varying risk profiles and dependencies on their endowments all seem to have similar betas ranging from 0.6 to 0.75. This situation may have its origin in the common historical starting point for most endowments and foundations, the conventional 60/40 stock and bond portfolio. For an understanding of what happens to these high-beta portfolios that contain a large percentage of alternative assets in extreme market conditions, the concept of stress beta introduced by Leibowitz,[12] is important and especially useful.

Risk Parity

Risk Parity is a special case of risk allocation. However, one may ask, "parity with respect to what?" In some cases it may be useful to think about a target volatility to compare different sources of return while building a portfolio. In the broader context, however, it may be more useful for an institution to develop the right risk sensitivities or allocation with regard to a series of possible economic and market scenarios that the portfolio or institution may be exposed to.

In the end the institution's investment strategy is simply part of a much broader effort to control the risk of not achieving institutional goals under a variety of possible market scenarios. The Wealth Allocation Framework clarifies the role of an investment strategy within the broader context of essential and aspirational goals and stable and unstable economic scenarios.

The Role of True Alpha

True alpha comes from idiosyncratic investing—tactical asset allocation, security selection and so on. However, for it to make a difference in the aggregate to the institution, this strategy must be accomplished on a large scale i.e. on a scale comparable to the institution's wealth. This is neither desirable nor feasible for most endowments and is accomplished only occasionally. Successful examples are few and far between and in most cases, it means charting one's own course and having the backing of the institution to stay the course for many decades[13].

As a rule, idiosyncratic return from superior investment and manager selection will only contribute a small percentage to the total wealth of an institution. Therefore, the risks taken to achieve these excess returns should be small in the aggregate. Not every investment portfolio can execute well on the Aspirational Portfolio. Instead, an institution may try to fill this bucket through other activities such as raising large donor gifts or identifying intellectual property that may yield potentially lucrative patents. Other examples may be developing a top-notch medical facility or a top-tier sports team. The overall portfolio construction should focus on achieving institutional goals. In the institutional sense, this demands a closer connection between the investment strategy and other activities of the institution.

Acknowledgements

I would like to thank Nevenka Vrdoljak and Ravindra Koneru for assistance with the figures and calculations for this paper. An earlier version of this paper was published in the NMS Exchange newsletter in November 2012.

Sources, References and Further Reading

[1] Harry Markowitz, The Journal of Finance, Vol. 7, No. 1. (1952)

[2] David F Swensen: Pioneering Portfolio Management, Published by Free Press (2000)

[3] William Goetzmann & Sharon Oster, Competition Among University Endowments, Yale ICF Working Paper No. 06-14, (2012)

[4] J. Tobin, Liquidity preference as behavior toward risk. Review of Economic Studies 25 (1958)

[5] Carmen M. Reinhart & Kenneth S. Rogoff, This Time Is Different: Eight Centuries of Financial Folly, Published by Princeton University Press (2011)

[6] Elroy Dimson, Paul Marsh, and Mike Staunton, Triumph of the Optimists: 101 Years of Global Investment Returns. Published by Princeton University Press (2002)

[7] B.B. Mandelbrot, The Random Character of Stock Market Prices ed. P Cootner, Published by MIT Press (1964)

[8] Nacubo-Commonfund Study 2013 (see *www.nacubo.org*)

[9] Ashvin B. Chhabra, Ravindra Koneru & Lex Zaharoff, “Modern Portfolio Theory’s Third Rail”, Journal of Wealth Management, Vol. 14, No. 1 (2011)

[10] Ashvin B. Chhabra, “Beyond Markowitz: A comprehensive Wealth Allocation Framework for individual investors”, Journal of Wealth Management, Vol. 7, No. 4 (2005)

[11] Ashvin B. Chhabra, “Clarifications on Beyond Markowitz”, Journal of Wealth Management, Vol. 11, No. 3 (2007)

[12] Martin L. Leibowitz, Anthony Bova, & P. Brett Hammond, The Endowment Model of Investing: Return, Risk, and Diversification, Published by Wiley Finance (2010)

[13] David Chambers and Elroy Dimson, Keynes the stock market investor, available at *http://ssrn.com/abstract=2023011*

GWM Investment Management & Guidance (IMG) provides industry-leading investment solutions, portfolio construction advice and wealth management guidance. This material was prepared by the Investment Management & Guidance Group (IMG) and is not a publication of BofA Merrill Lynch Global Research. The views expressed are those of IMG only and are subject to change. This information should not be construed as investment advice. It is presented for information purposes only and is not intended to be either a specific offer by any Merrill Lynch entity to sell or provide, or a specific invitation for a consumer to apply for, any particular retail financial product or service that may be available.

This information and any discussion should not be construed as a personalized and individual client recommendation, which should be based on each client's investment objectives, risk tolerance, liquidity needs and financial situation. This information and any discussion also is not intended as a specific offer by Merrill Lynch, its affiliates, or any related entity to sell or provide, or a specific invitation for a consumer to apply for, any particular retail financial product or service. Investments and opinions are subject to change due to market conditions and the opinions and guidance may not be profitable or realized. Any information presented in connection with BofA Merrill Lynch Global Research is general in nature and is not intended to provide personal investment advice. The information does not take into account the specific investment objectives, financial situation and particular needs of any specific person who may receive it. Investors should understand that statements regarding future prospects may not be realized.

Asset allocation and diversification do not assure a profit or protect against a loss during declining markets.

© 2014 Bank of America Corporation. All rights reserved.

Bank of America
Merrill Lynch