AN INTENTION-BASED SEMANTICS FOR IMPERATIVES

Daniel W. Harris
Hunter College, CUNY

(1) **Do the right thing**
(2) **Snow is white**
(3) **Fly me to the moon** and **let me play among the stars**
(4) **Make us omelettes** or **I'll get us some bagels**
(5) **Help me** if **you can**

$[\![\text{Snow is white}]\!]$ = ???

DECLARATIVE

$[\![\text{Do the right thing}]\!]$ = ???

IMPERATIVE

TWO ASSUMPTIONS

TWO ASSUMPTIONS

1. Clauses factor, at LF, into a mood-marker and a moodless sentence radical.

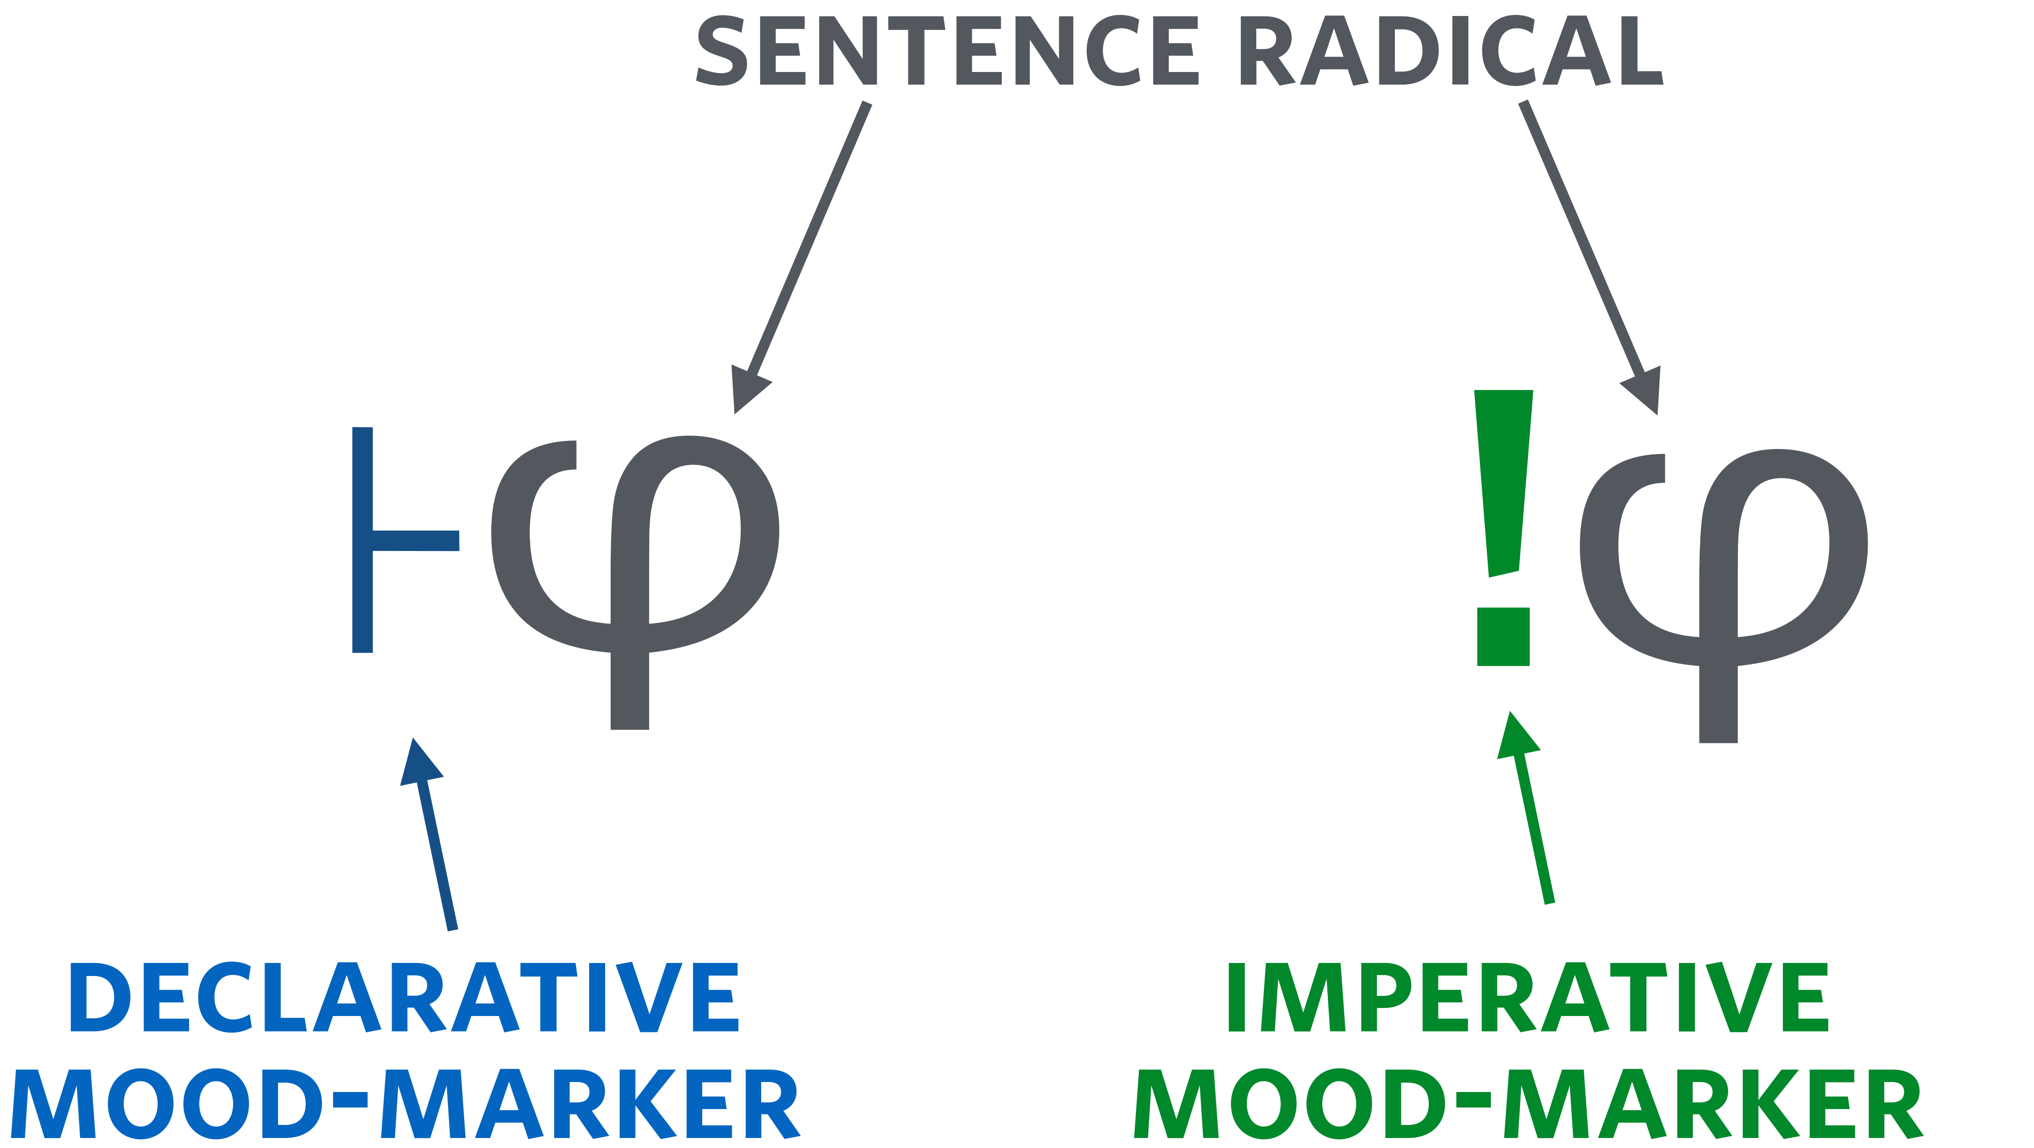

TWO ASSUMPTIONS

2. The semantic value of a sentence radical is a set of possible worlds.

$$(\forall \varphi)\, [\![\varphi]\!] \in W$$

POSITIVE VIEW: A SKETCH

cf. Charlow (2014)

POSITIVE VIEW: A SKETCH

For any sentence radical φ:

$[\![\vdash\varphi]\!]^c$ is a belief

(namely: the belief that $[\![\varphi]\!]^c$ is true)

$[\![!\varphi]\!]^c$ is an intention

(namely: the intention to make $[\![\varphi]\!]^c$ true)

POSITIVE VIEW: A SKETCH

$$[\![\text{snow is white}]\!]^c =$$

The belief that snow is white.

POSITIVE VIEW: A SKETCH

$$[\![\text{buy me a drink}]\!]^c =$$

the intention to buy speaker_c a drink

POSITIVE VIEW: A SKETCH

How to
formalize this?

POSITIVE VIEW: A SKETCH

M_A

A's **COGNITIVE MODEL**
=a set-theoretic representation of A's beliefs and plans

M_A determines:
(in a way to be explained)

B_A

A's **BELIEF STATE**
=the set of worlds compatible with what A believes

I_A

A's **INTENTION STATE**
=the set of worlds compatible with what A intends

POSITIVE VIEW: A SKETCH

Beliefs and intentions are formalized as properties of cognitive models:

Belief that dogs are better than cats:

$\lambda \mathbf{M_A} \,.\, \mathbf{B_A} \subseteq \{w : \text{dogs are better than cats at } w\}$

Intention to high-five Beyoncé:

$\lambda \mathbf{M_A} \,.\, \mathbf{I_A} \subseteq \{w : \text{A high-fives Beyonce at } w\}$

POSITIVE VIEW: A SKETCH

So, the semantic values of clauses are properties of cognitive models too.

For any sentence radical φ,

$$[\![\vdash\varphi]\!]^c = \lambda \mathbf{M}_{\mathbf{A}} \,.\, \mathbf{B}_{\mathbf{A}} \subseteq [\![\varphi]\!]^c$$

$$[\![!\varphi]\!]^c = \lambda \mathbf{M}_{\mathbf{A}} \,.\, \mathbf{I}_{\mathbf{A}} \subseteq [\![\varphi]\!]^c$$

ALTERNATIVE TREATMENTS

FIRST ALTERNATIVE:
STATIC SEMANTICS, DYNAMIC PRAGMATICS

(Portner; von Fintel & Iatradou; Roberts)

$[\![\text{Snow is white.}]\!] =$

$\lambda w_{st} .$ snow is white in w

(The proposition that snow is white.)

$[\![\text{Do the right thing!}]\!] =$

$\lambda w_{st} . \lambda x_e : x = \alpha_c .$ x does the right thing in w

(A property (restricted to the addressee) of doing the right thing.)

PROBLEM FOR STATIC VIEWS:
MIXED COORDINATION

(We're about to go into the bar together. I say:)

Buy us drinks and I'll find a table.

NOTE:

- Needn't have a conditional meaning.
- Can mean roughly: 'I'll find a table. Buy me a drink.'
- Can be the consequent of a conditional:
  'If your friend is tending bar, buy us drinks and I'll find a table.'

PROBLEM FOR STATIC VIEWS:
MIXED COORDINATION

(We're at a book store. Each of us has three books, but we only have enough money for five, total:)

Put back Naked Lunch or I'll put back Waverley.

(Starr ms)

NOTE:

- Needn't have a conditional meaning.
- Can be the consequent of a conditional:
  'If we only have $5, put back Naked Lunch or I'll put back Waverly.'

PROBLEM FOR STATIC VIEWS:
IMPERATIVE INFERENCE

Buy me a drink.

You won't buy me a drink unless you go to the bar.

↳ So, go to the bar!

Attack if the weather is good.

The weather is good.

↳ So, attack!

SECOND ALTERNATIVE:
DYNAMIC SEMANTICS

(e.g., Starr)

Starr's clauses:

$$R[\vdash\varphi] = \{r_\varphi \mid r \in R \;\&\; r_\varphi \neq \varnothing\}$$

(where $r_\varphi = \{\langle a[\varphi], a'[\varphi]\rangle \mid \langle a, a'\rangle \in r \;\&\; a[\varphi] \neq \varnothing\}$)

(A CCP that adds the content of **φ** to the context's information.)

$$R[!\varphi] = \{r \cup \{\langle c_r[\varphi], c_r - c_r[\varphi]\rangle\} \mid r \in R\}$$

(A CCP that adds a preference for $[\![\boldsymbol{\varphi}]\!]^c$ over $[\![\text{not } \boldsymbol{\varphi}]\!]^c$ to the context)

FIRST PROBLEM FOR DYNAMIC VIEWS:
RECALCITRANT SPEECH ACTS

“Stop what you’re doing.”

“We’ve been all wrong about this.”

FIRST PROBLEM FOR DYNAMIC VIEWS:

RECALCITRANT SPEECH ACTS

"…it would be a mistake to augment the theory of assertion with the fruits of the epistemological literature on belief-revision."

—N. Charlow (2014: §5.6.2)

FIRST PROBLEM FOR DYNAMIC VIEWS:
RECALCITRANT SPEECH ACTS

Solution:
Semantics specifies properties of cognitive models, but leaves it up to a theory of non-monotonic reasoning to sort out how addressees should satisfy those properties on particular occasions.

SECOND PROBLEM FOR DYNAMIC VIEWS:

BAD THEORY OF SPEECH ACTS

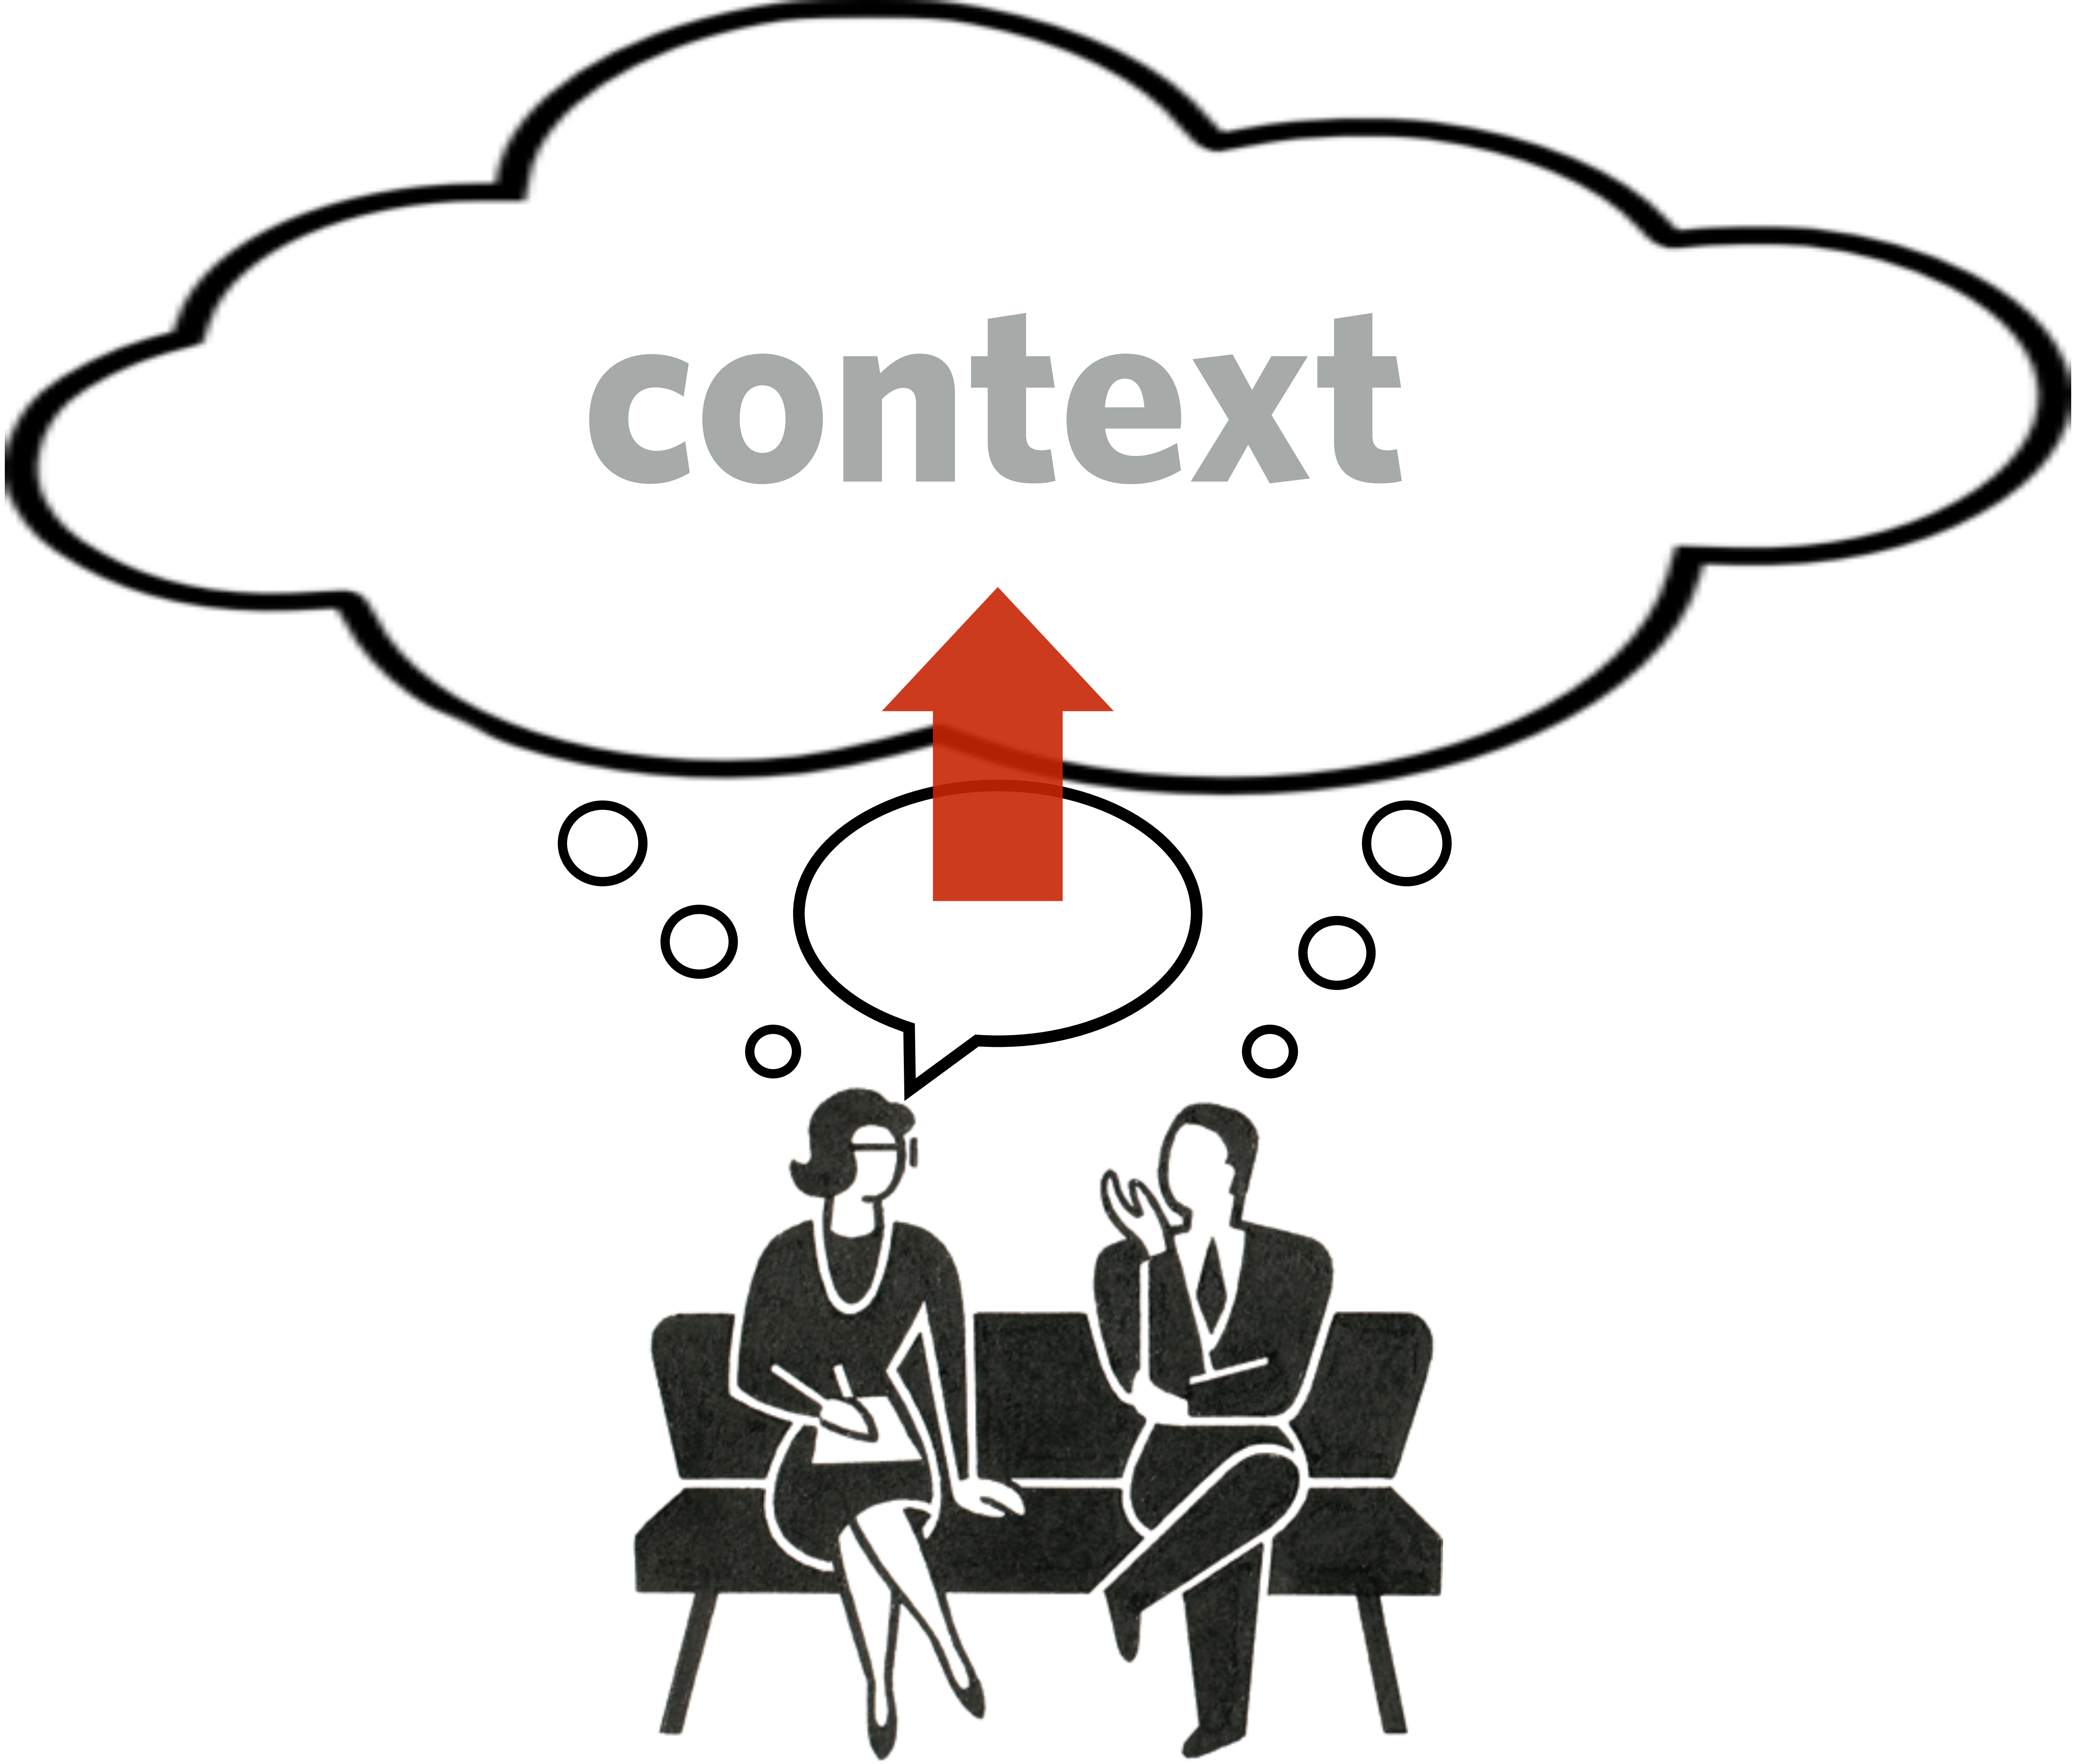
context

STALNAKER (1978, 2014):

To **assert** ***q*** is to propose adding *p* to the Common Ground (CG).

ROBERTS (1996/2012):

To **ask** ***q*** is to proffer ***q***, intending that it be adopted as the new Question Under Discussion (QUD).

PORTNER (2004):

To **direct A to φ** is to propose that **φ** be added to their section of the conversation's To-Do List (TDL).

CONTEXTS ARE PUBLIC

A proposition *p* is common ground of a conversation iff the participants *commonly accept p*:

- each accepts *p*;
- each accepts that each accepts *p*;
- etc.

(Stalnaker 2014)

A PRAGMATIC ARGUMENT

Roughly: We regularly perform speech acts and successfully communicate, in situations where we can't, and can't expect to, change the common ground.


S

The Coordinated Attack Problem

(The Byzantine Generals Problem)

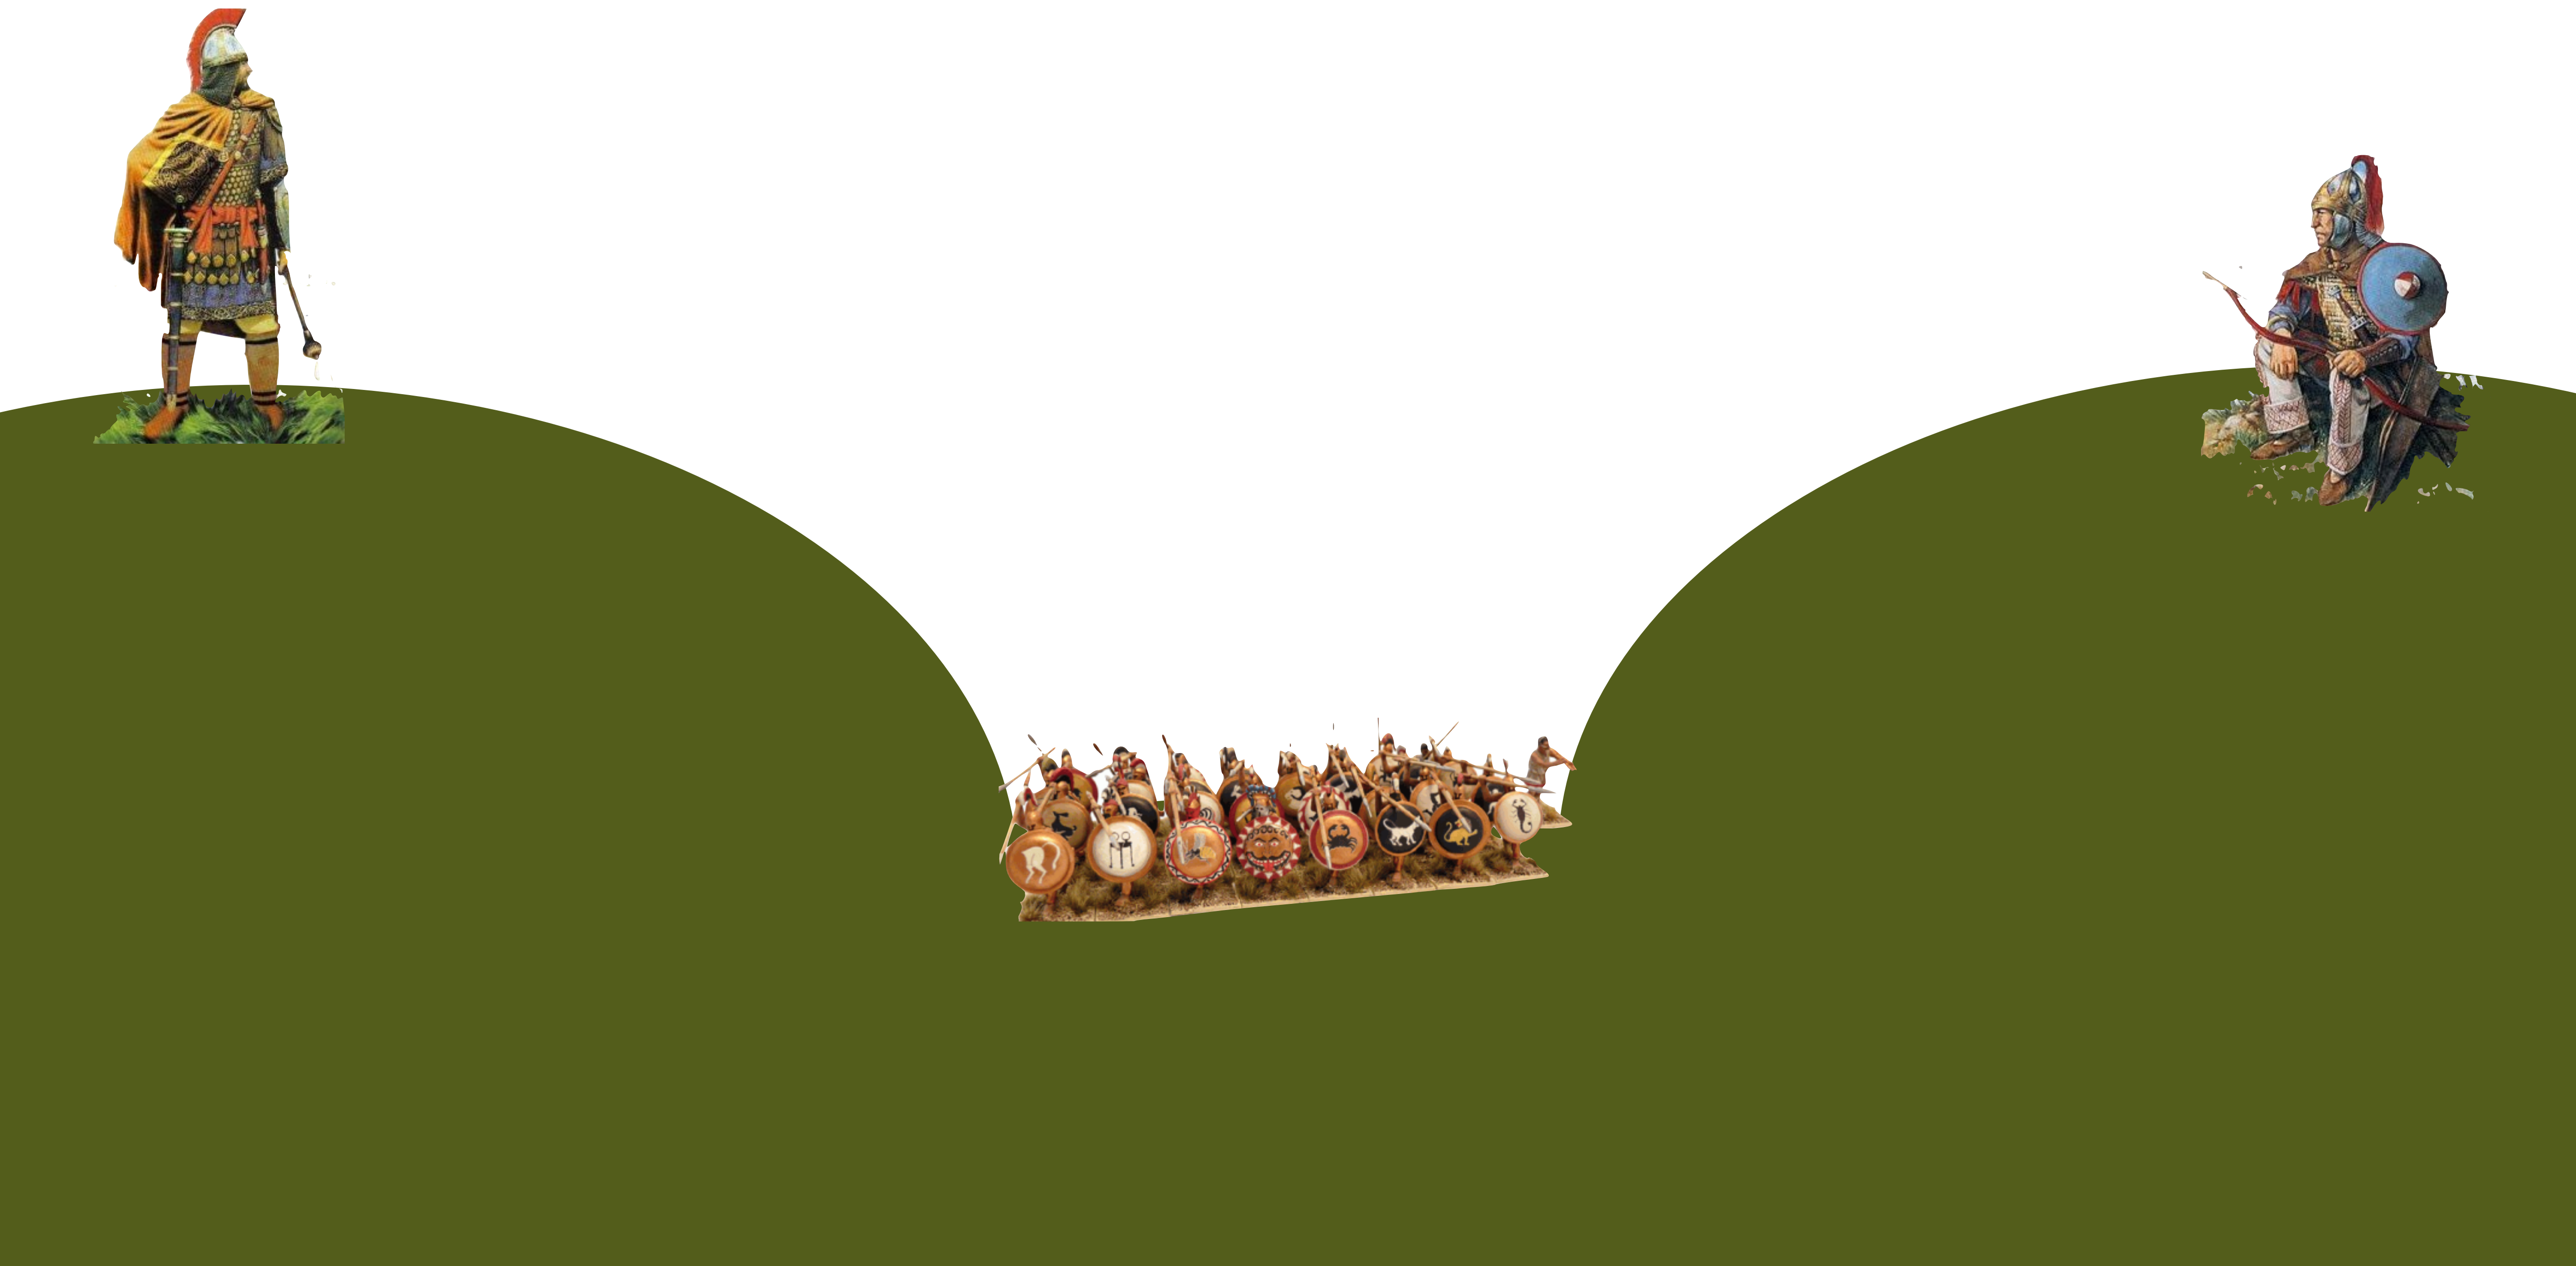

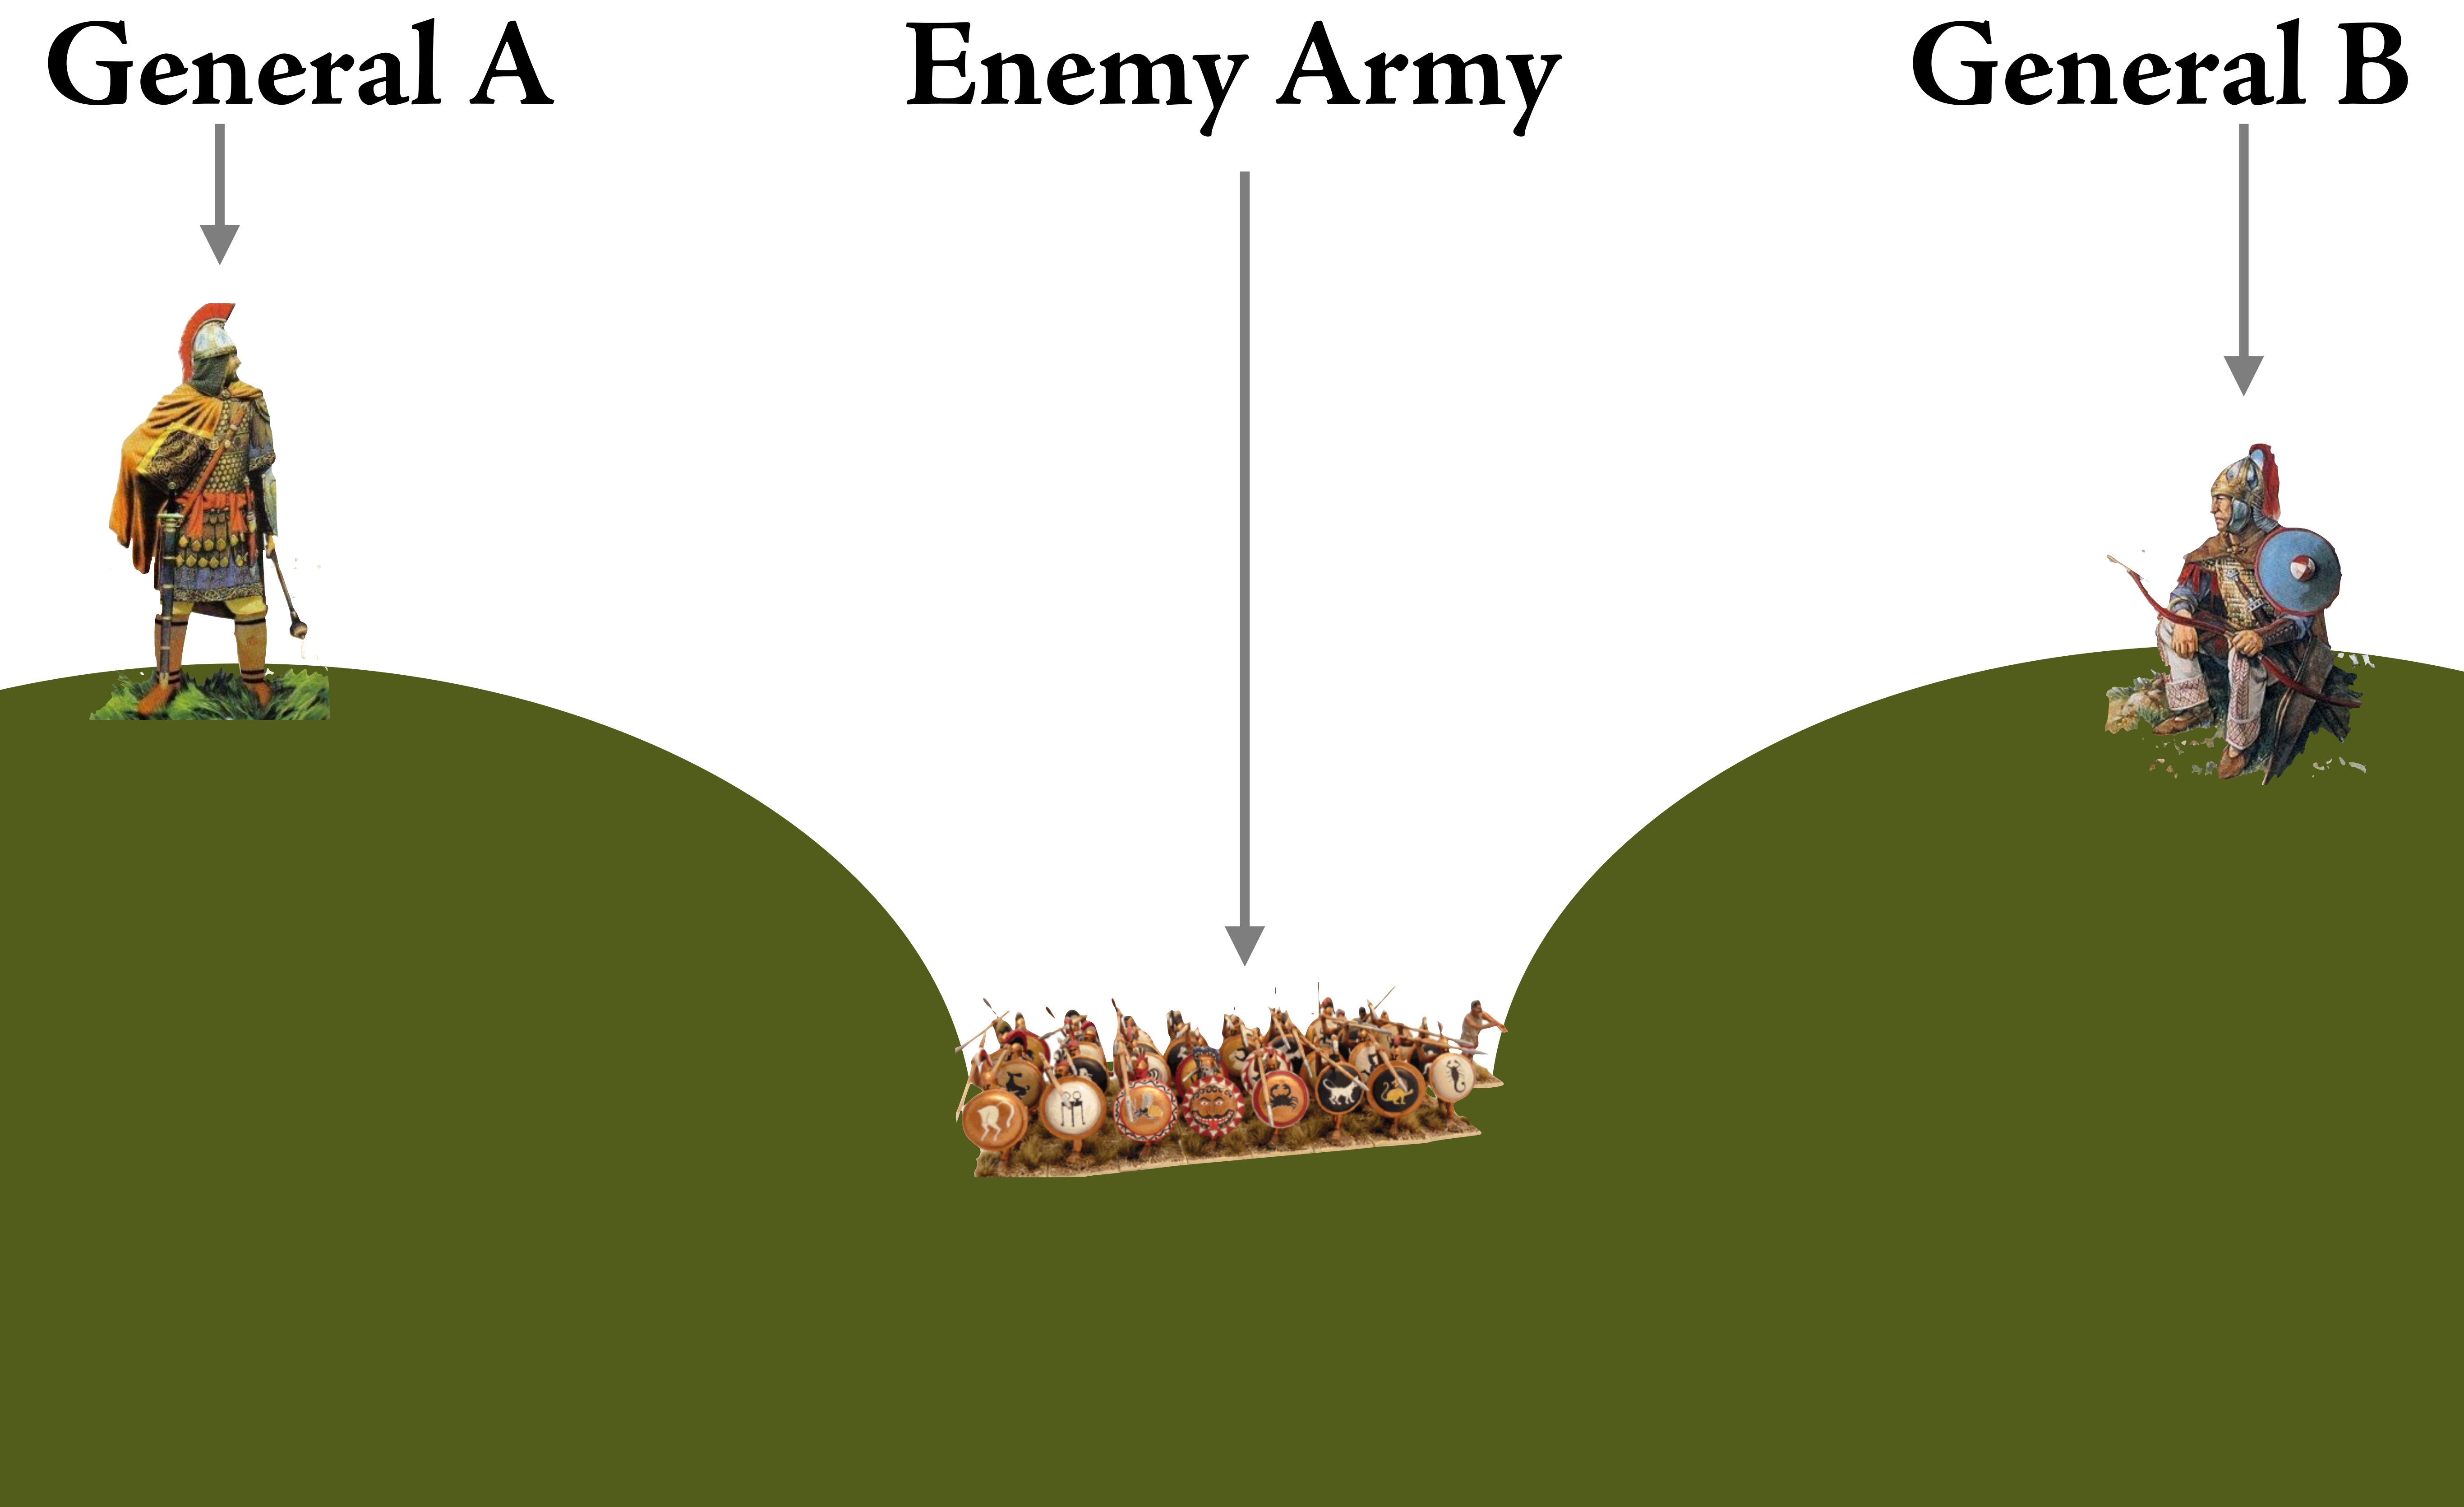
General A
Enemy Army
General B

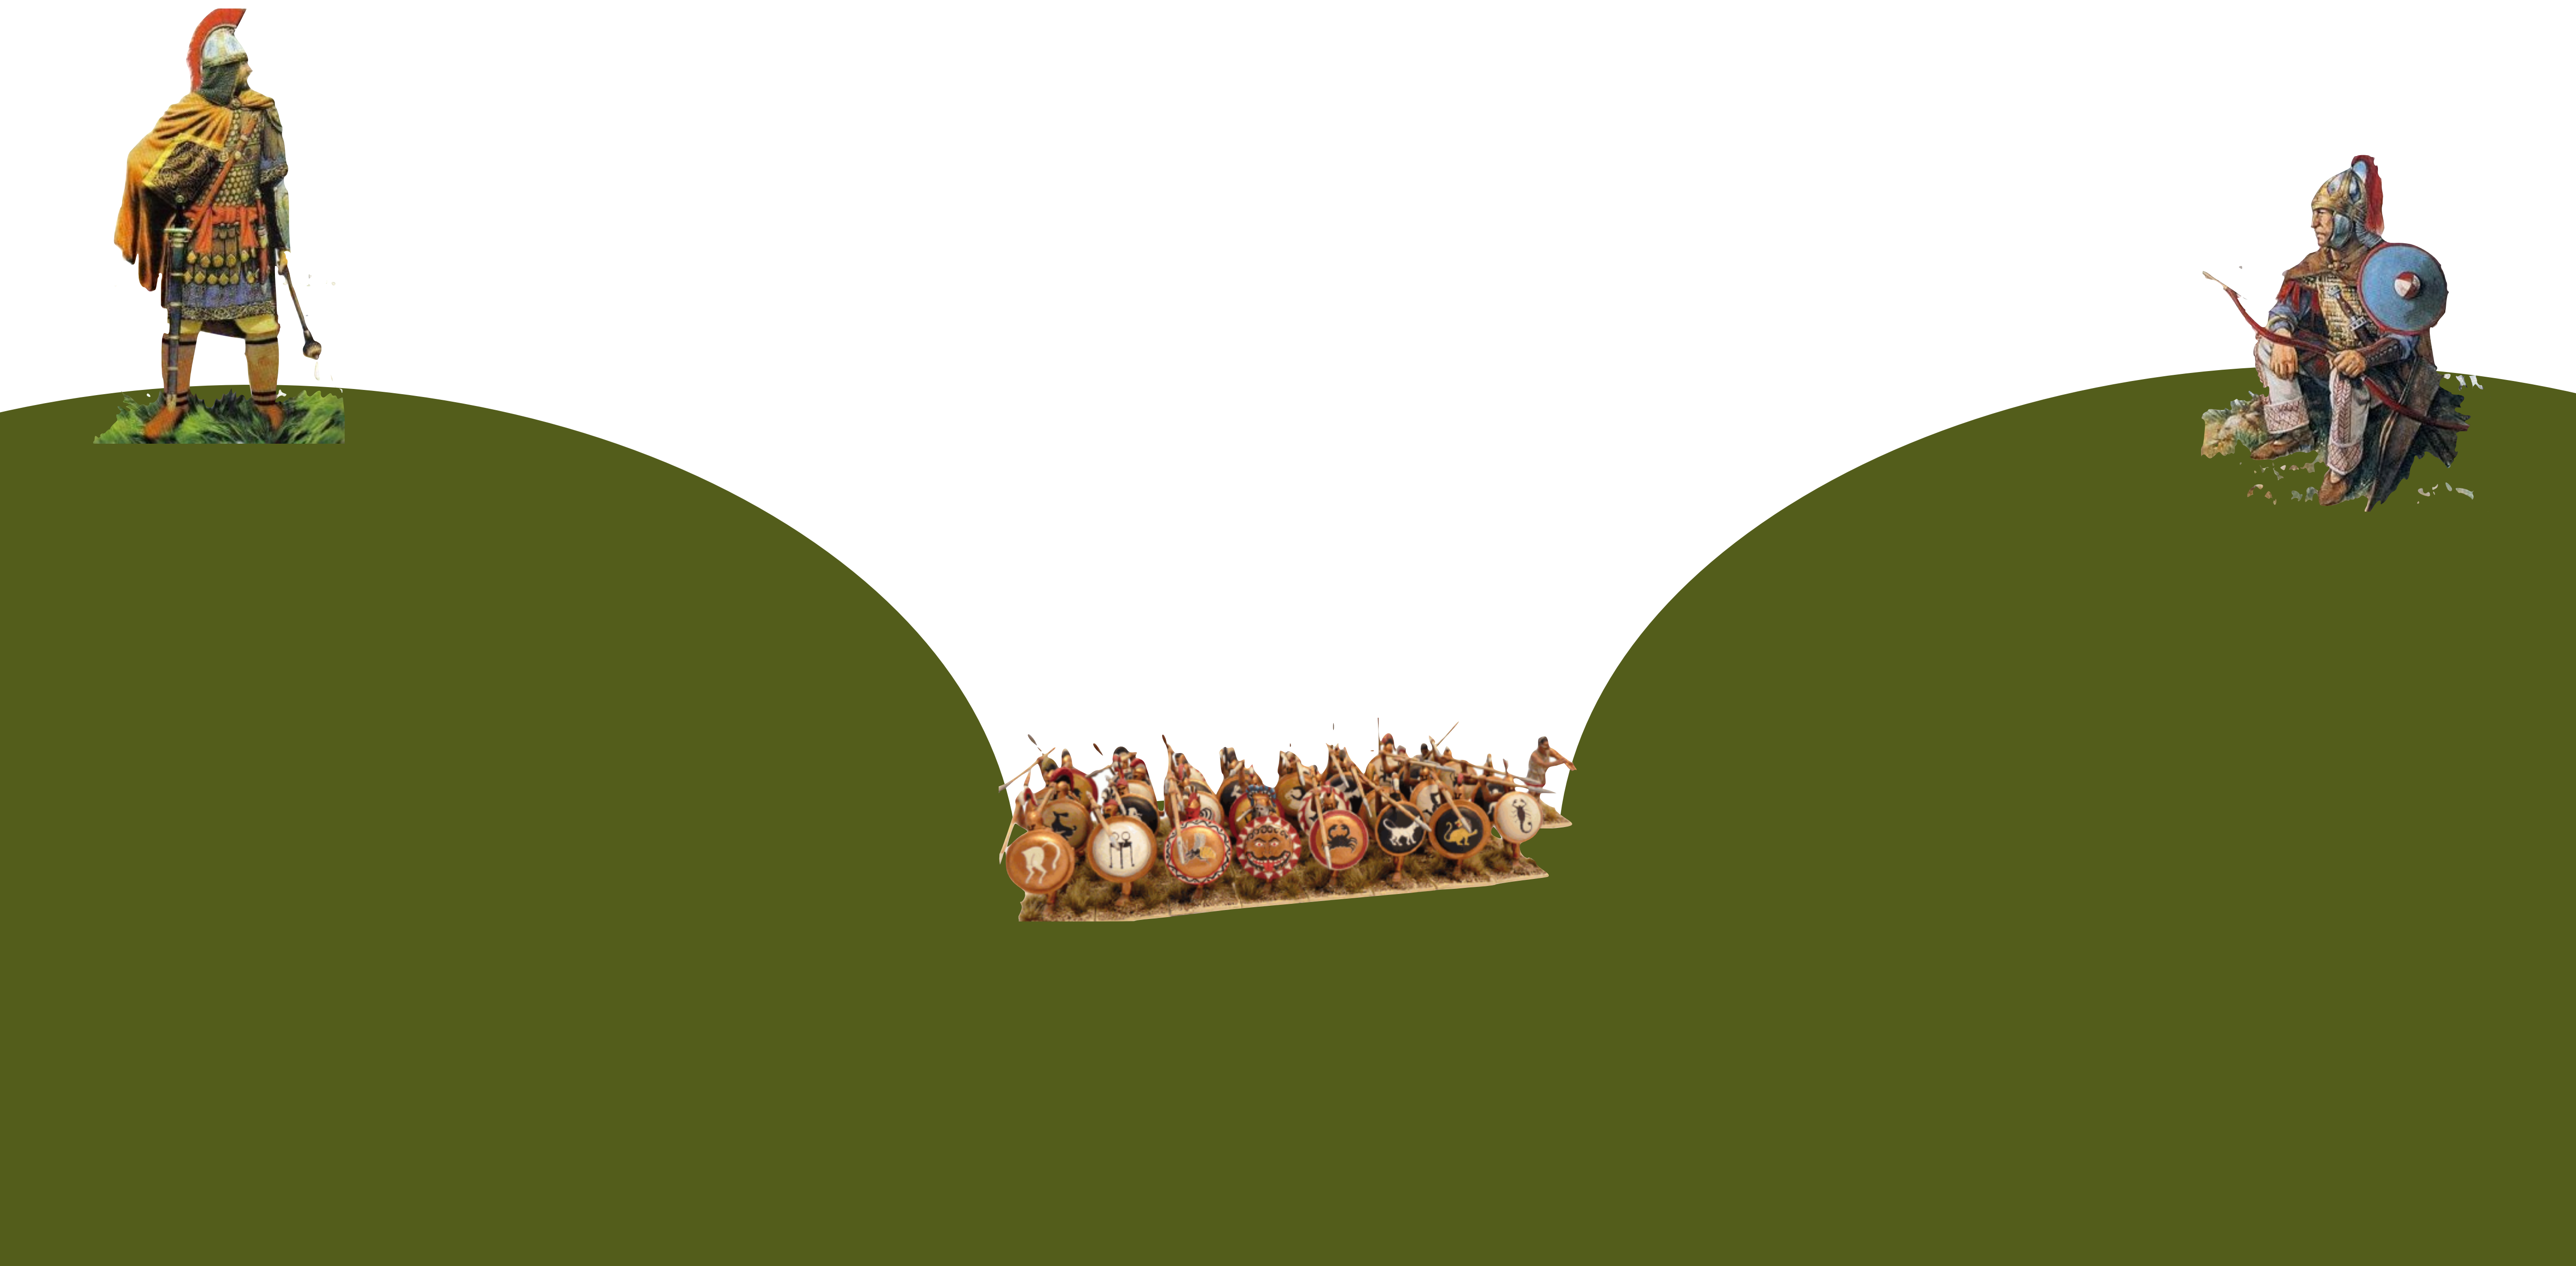

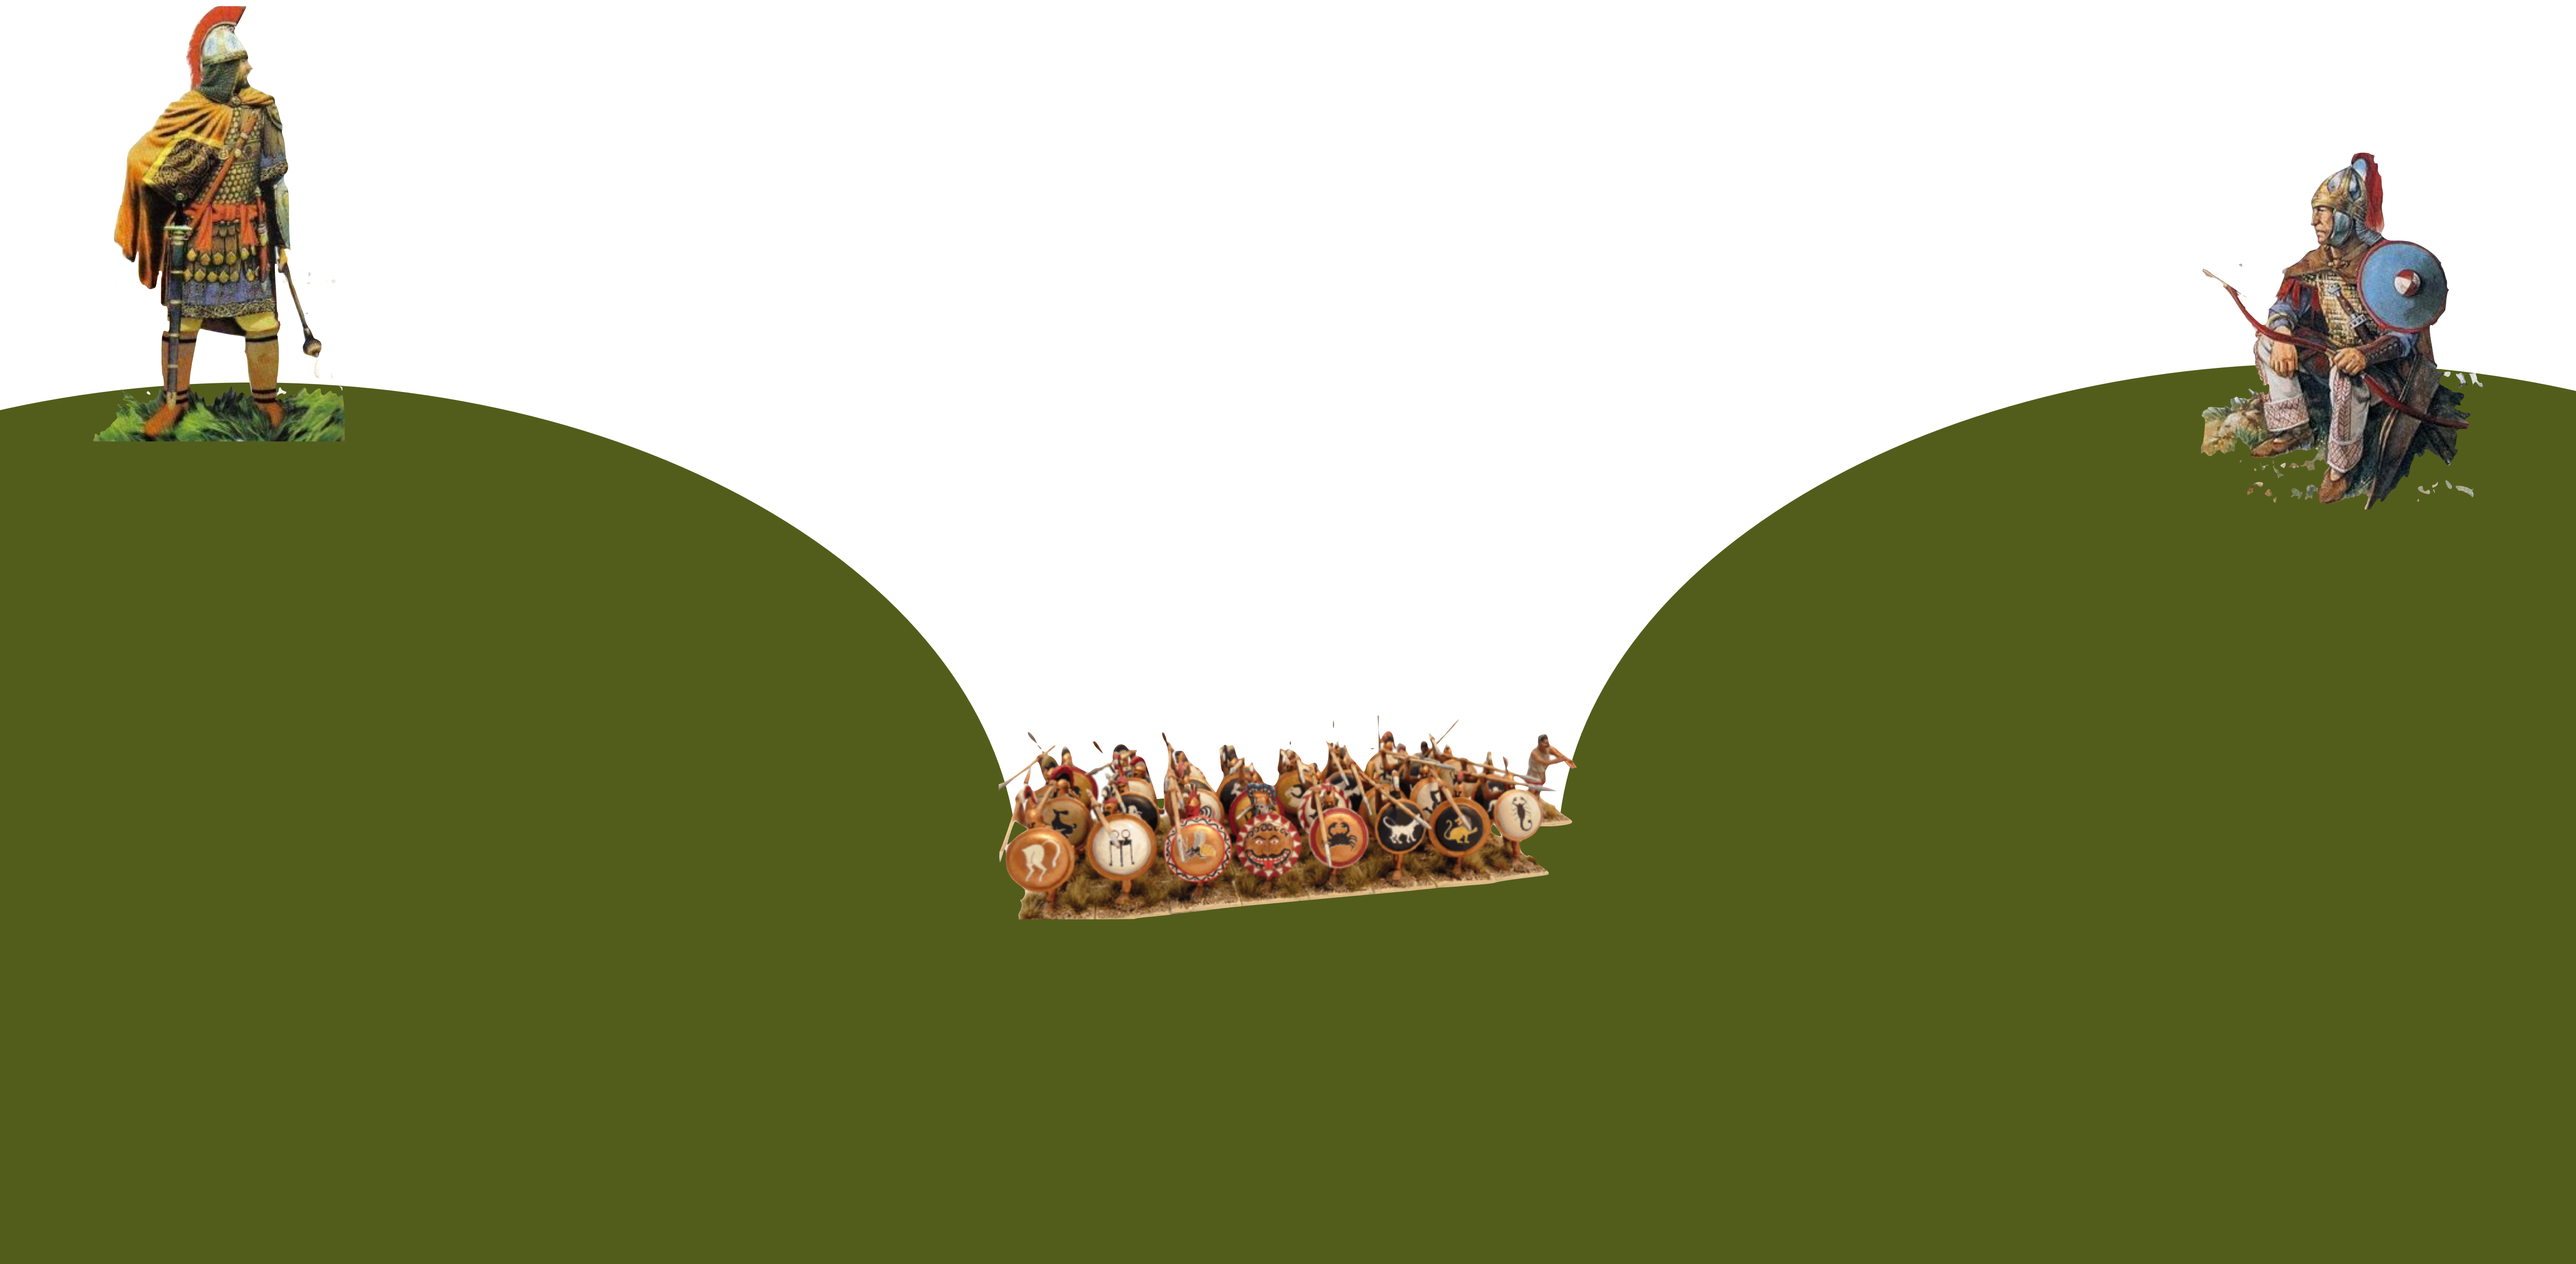

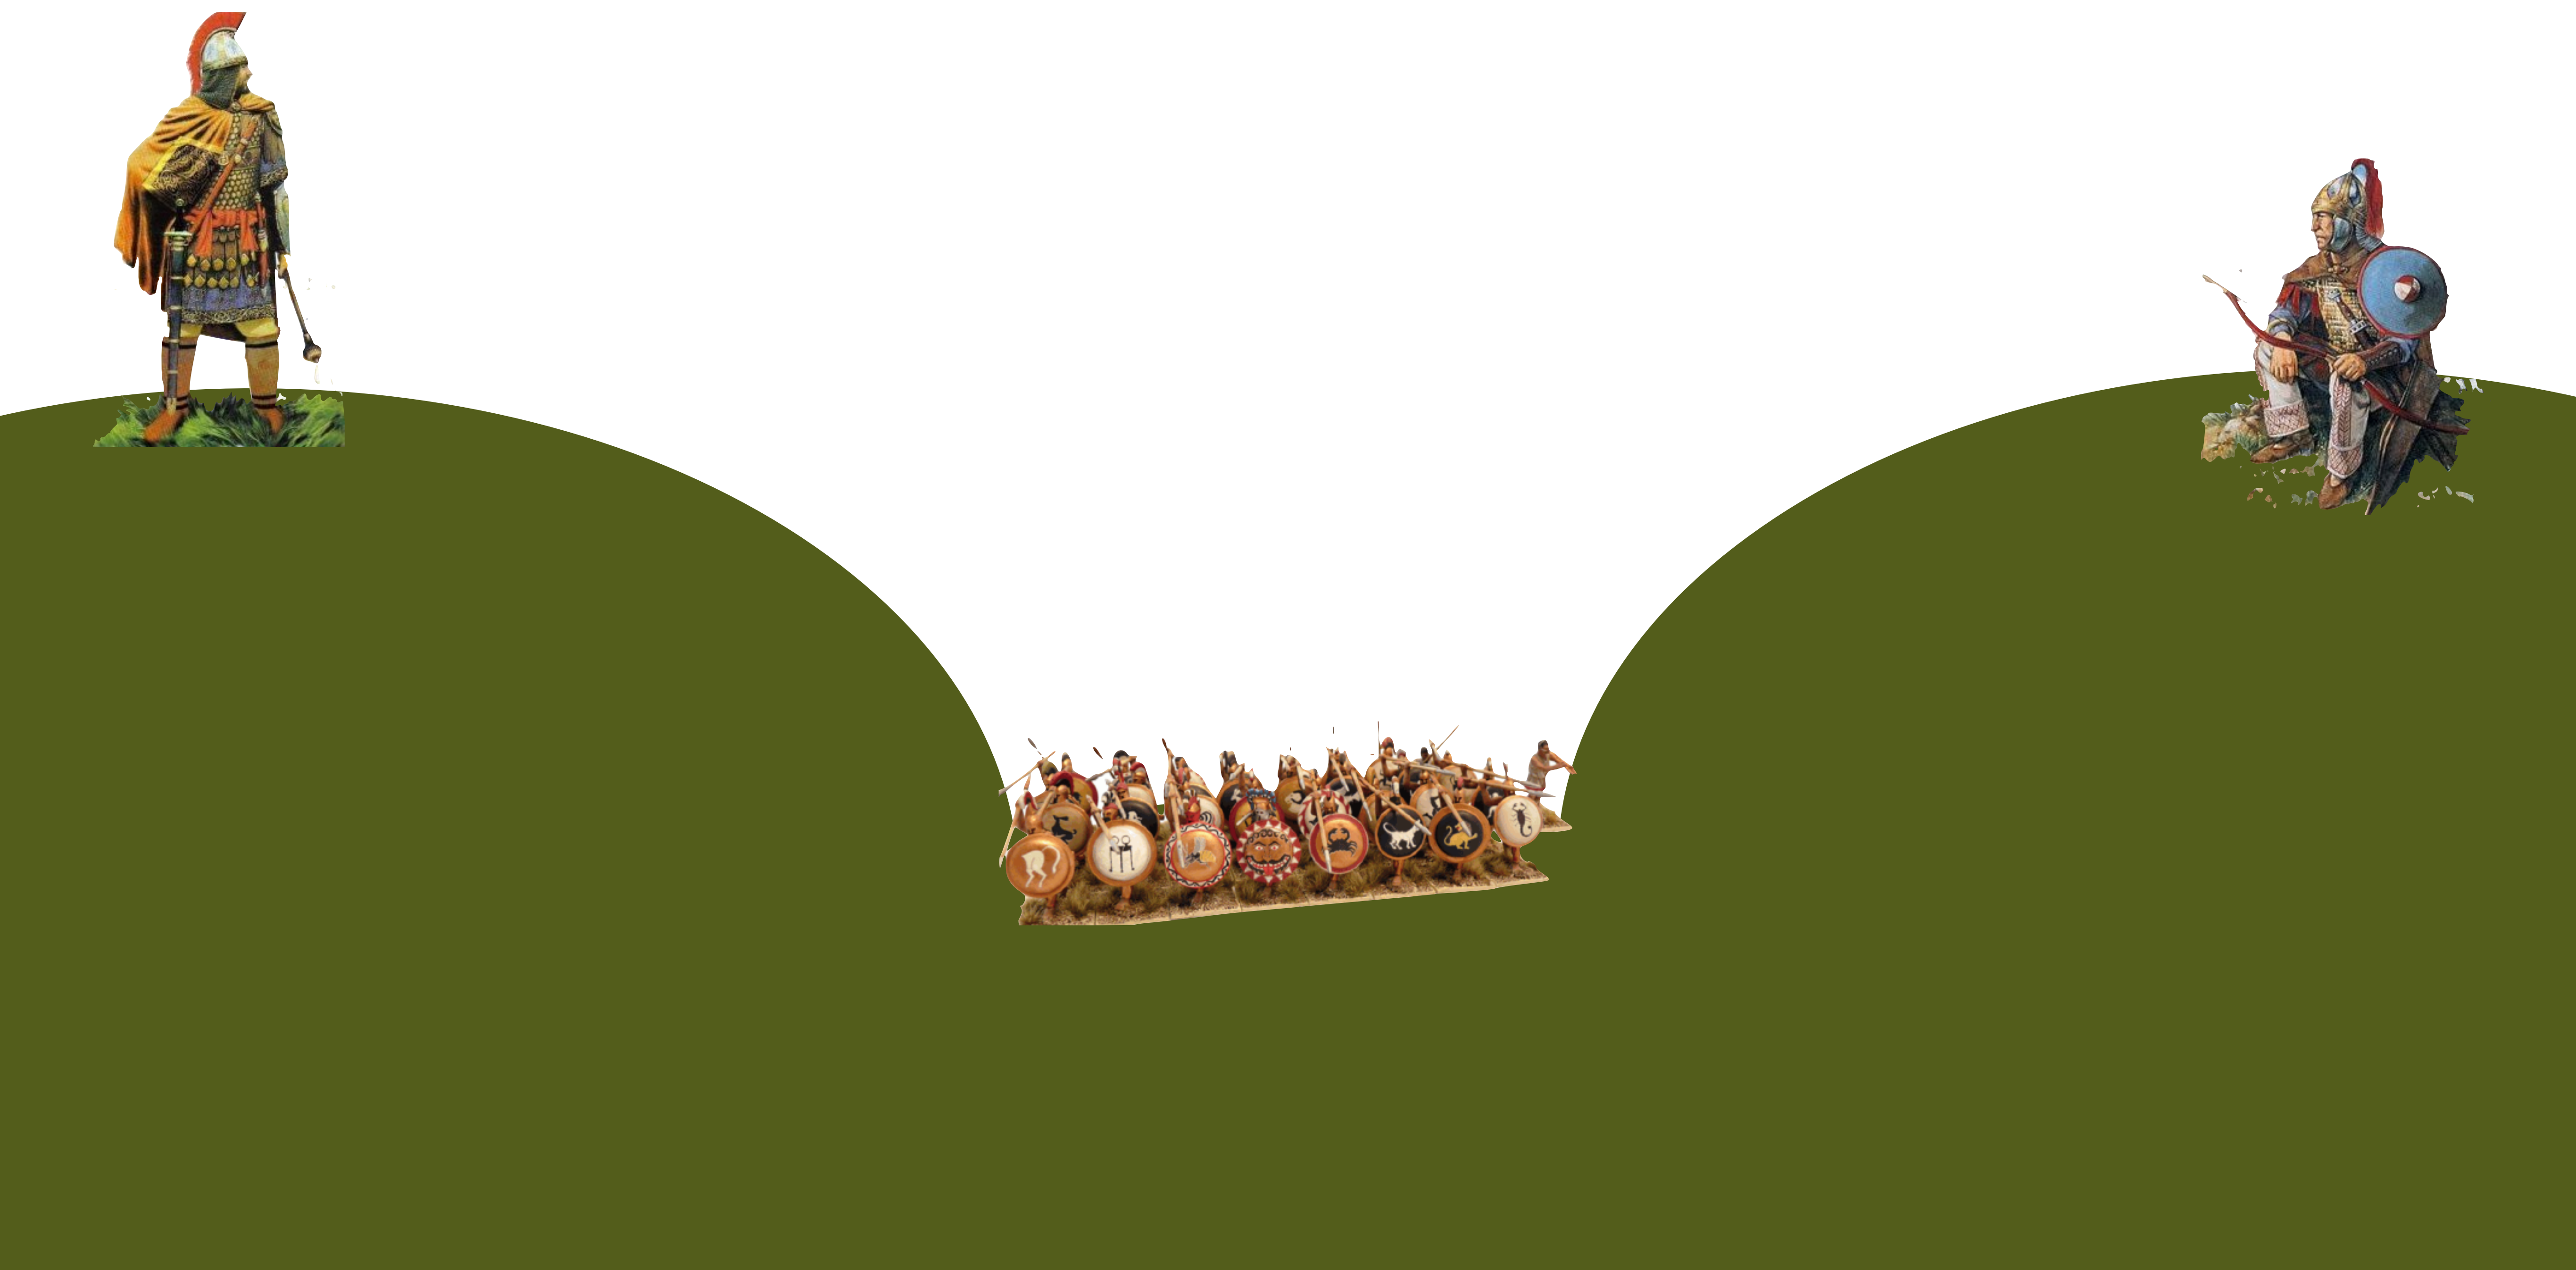

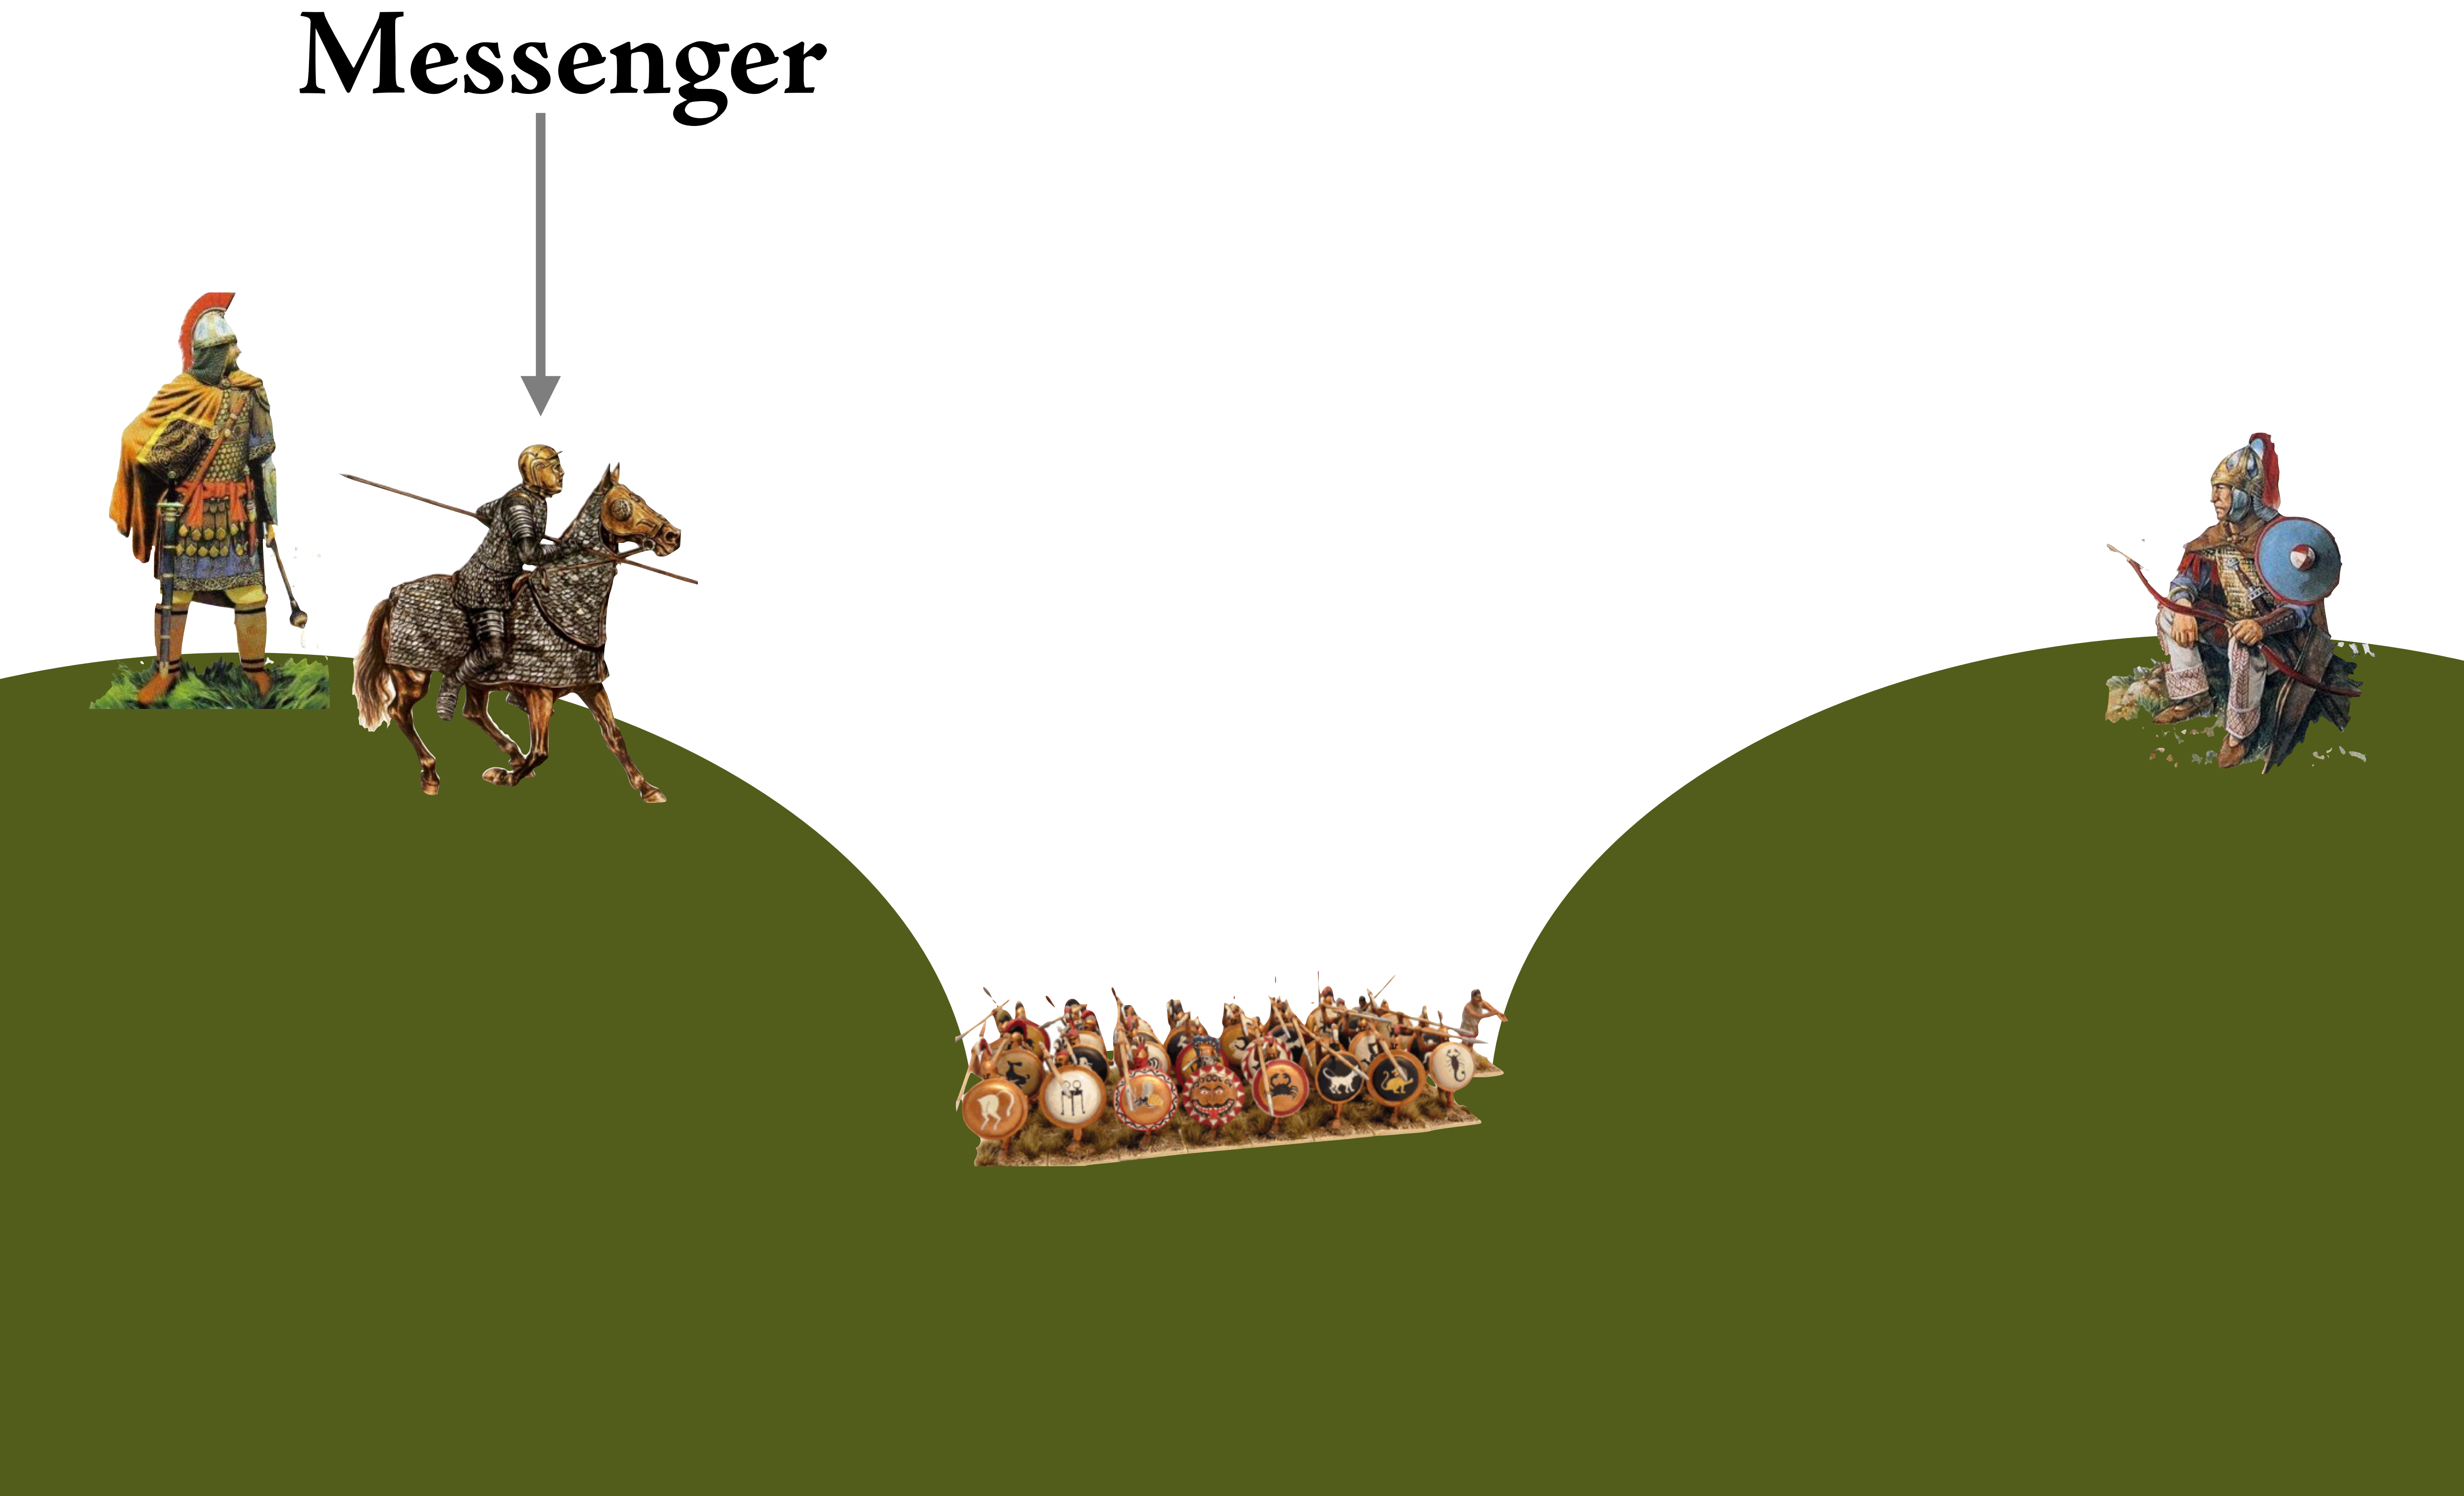
Messenger

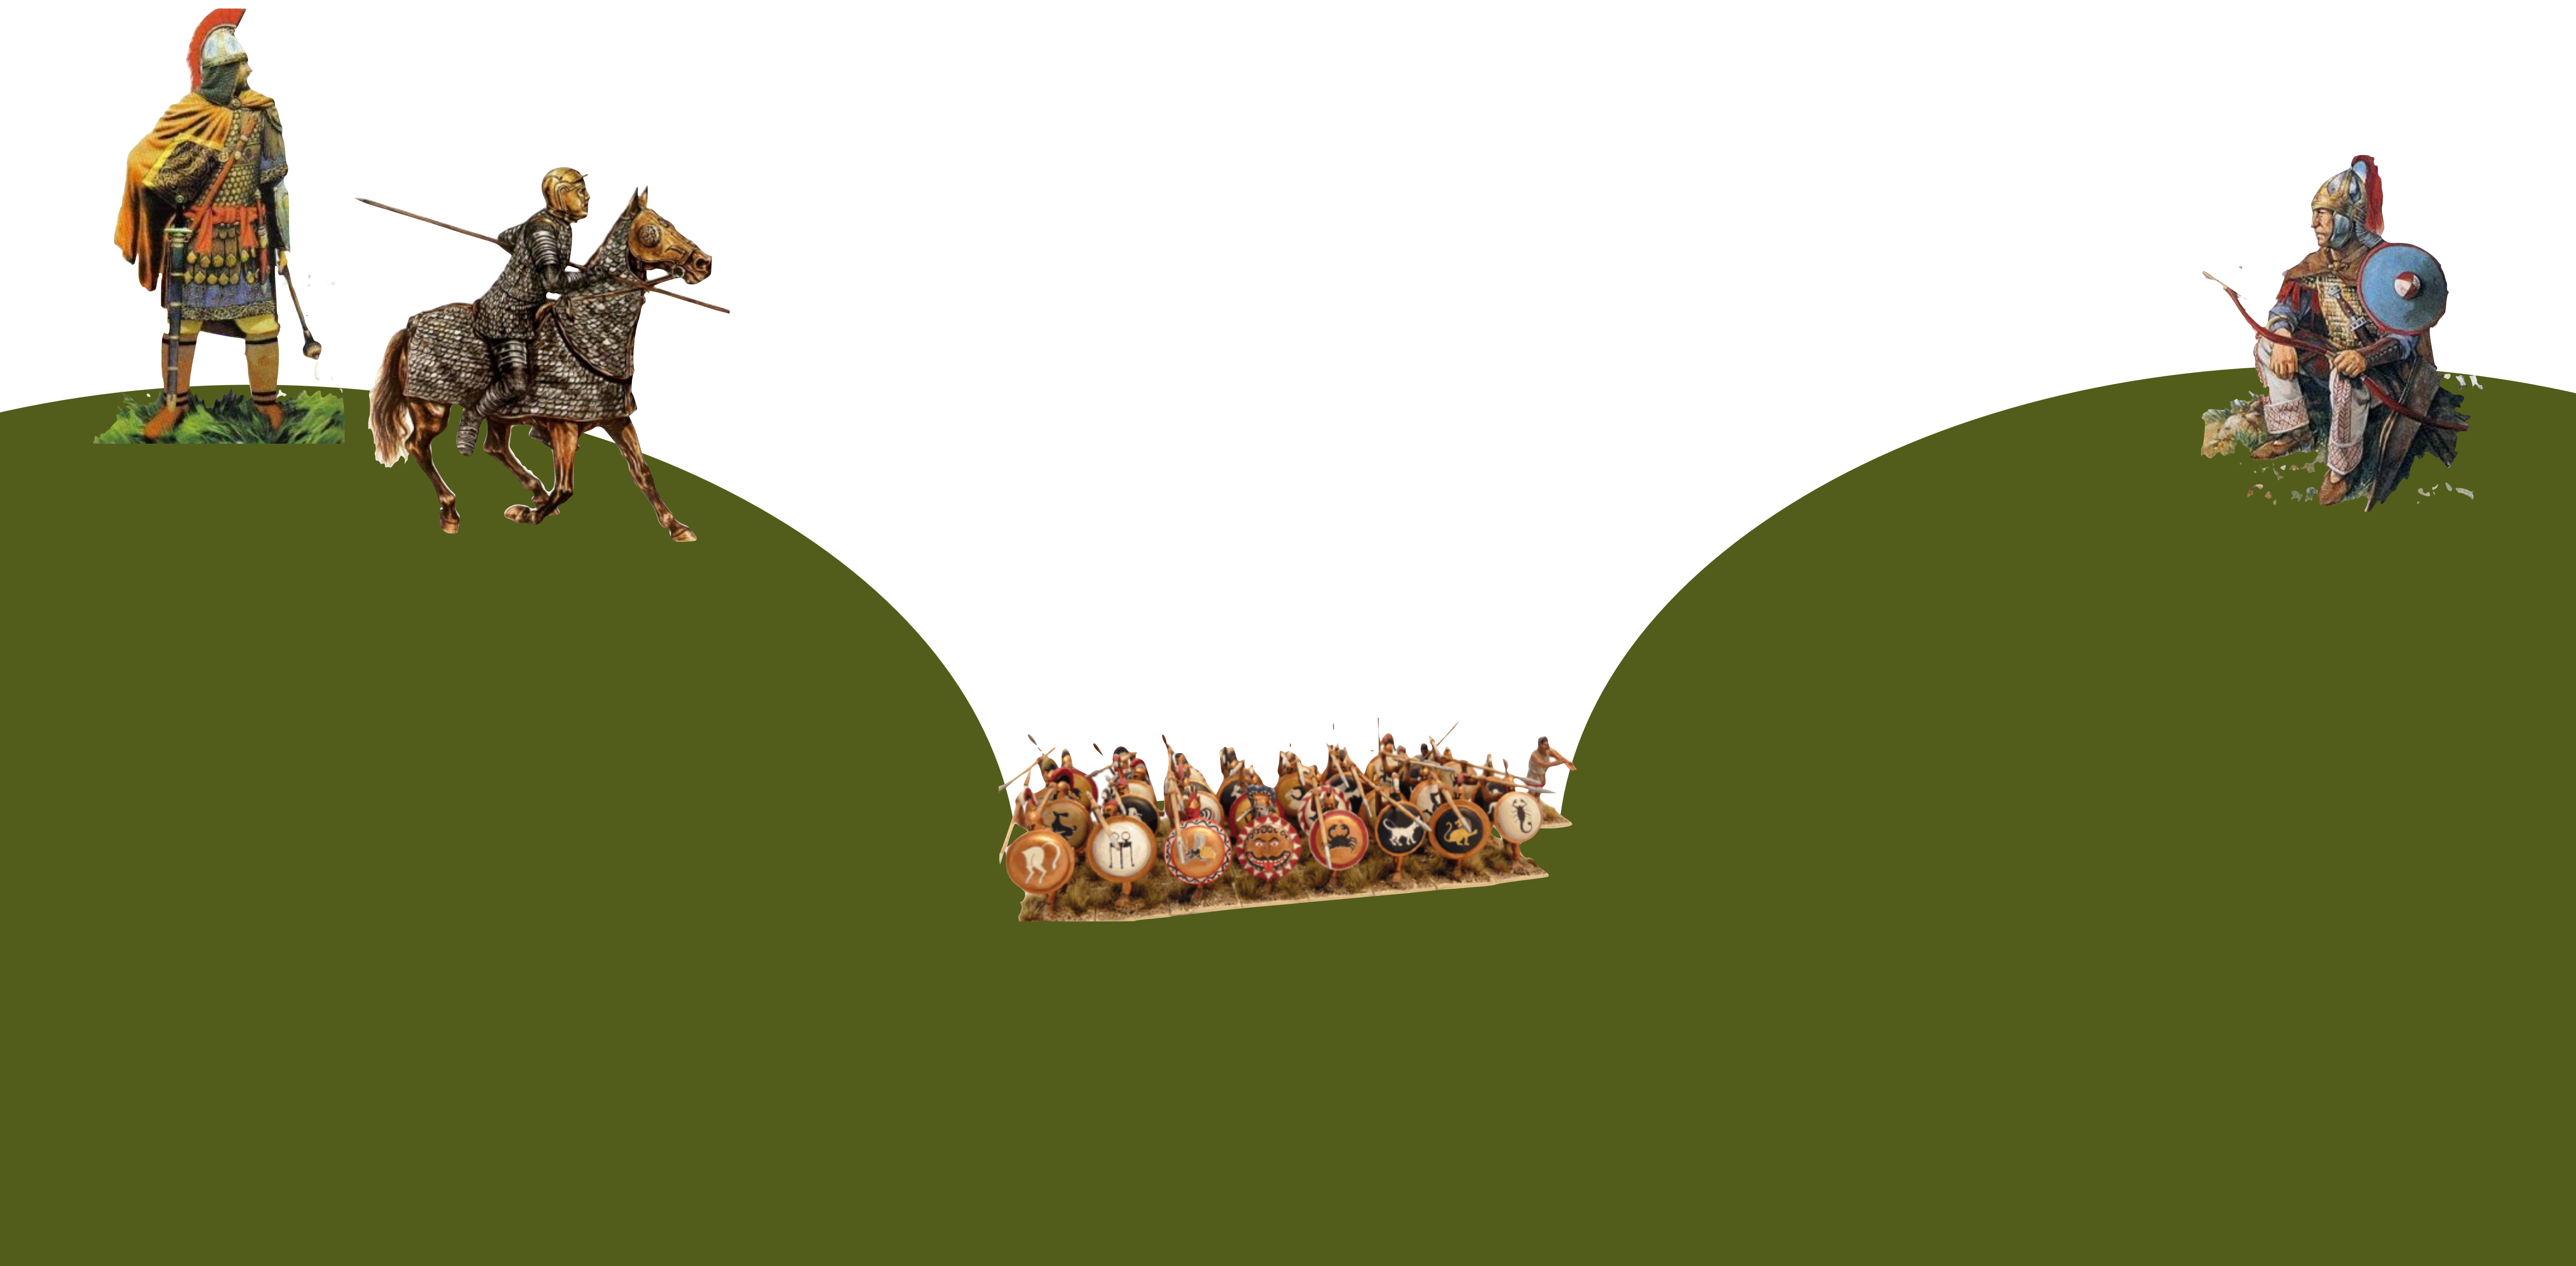

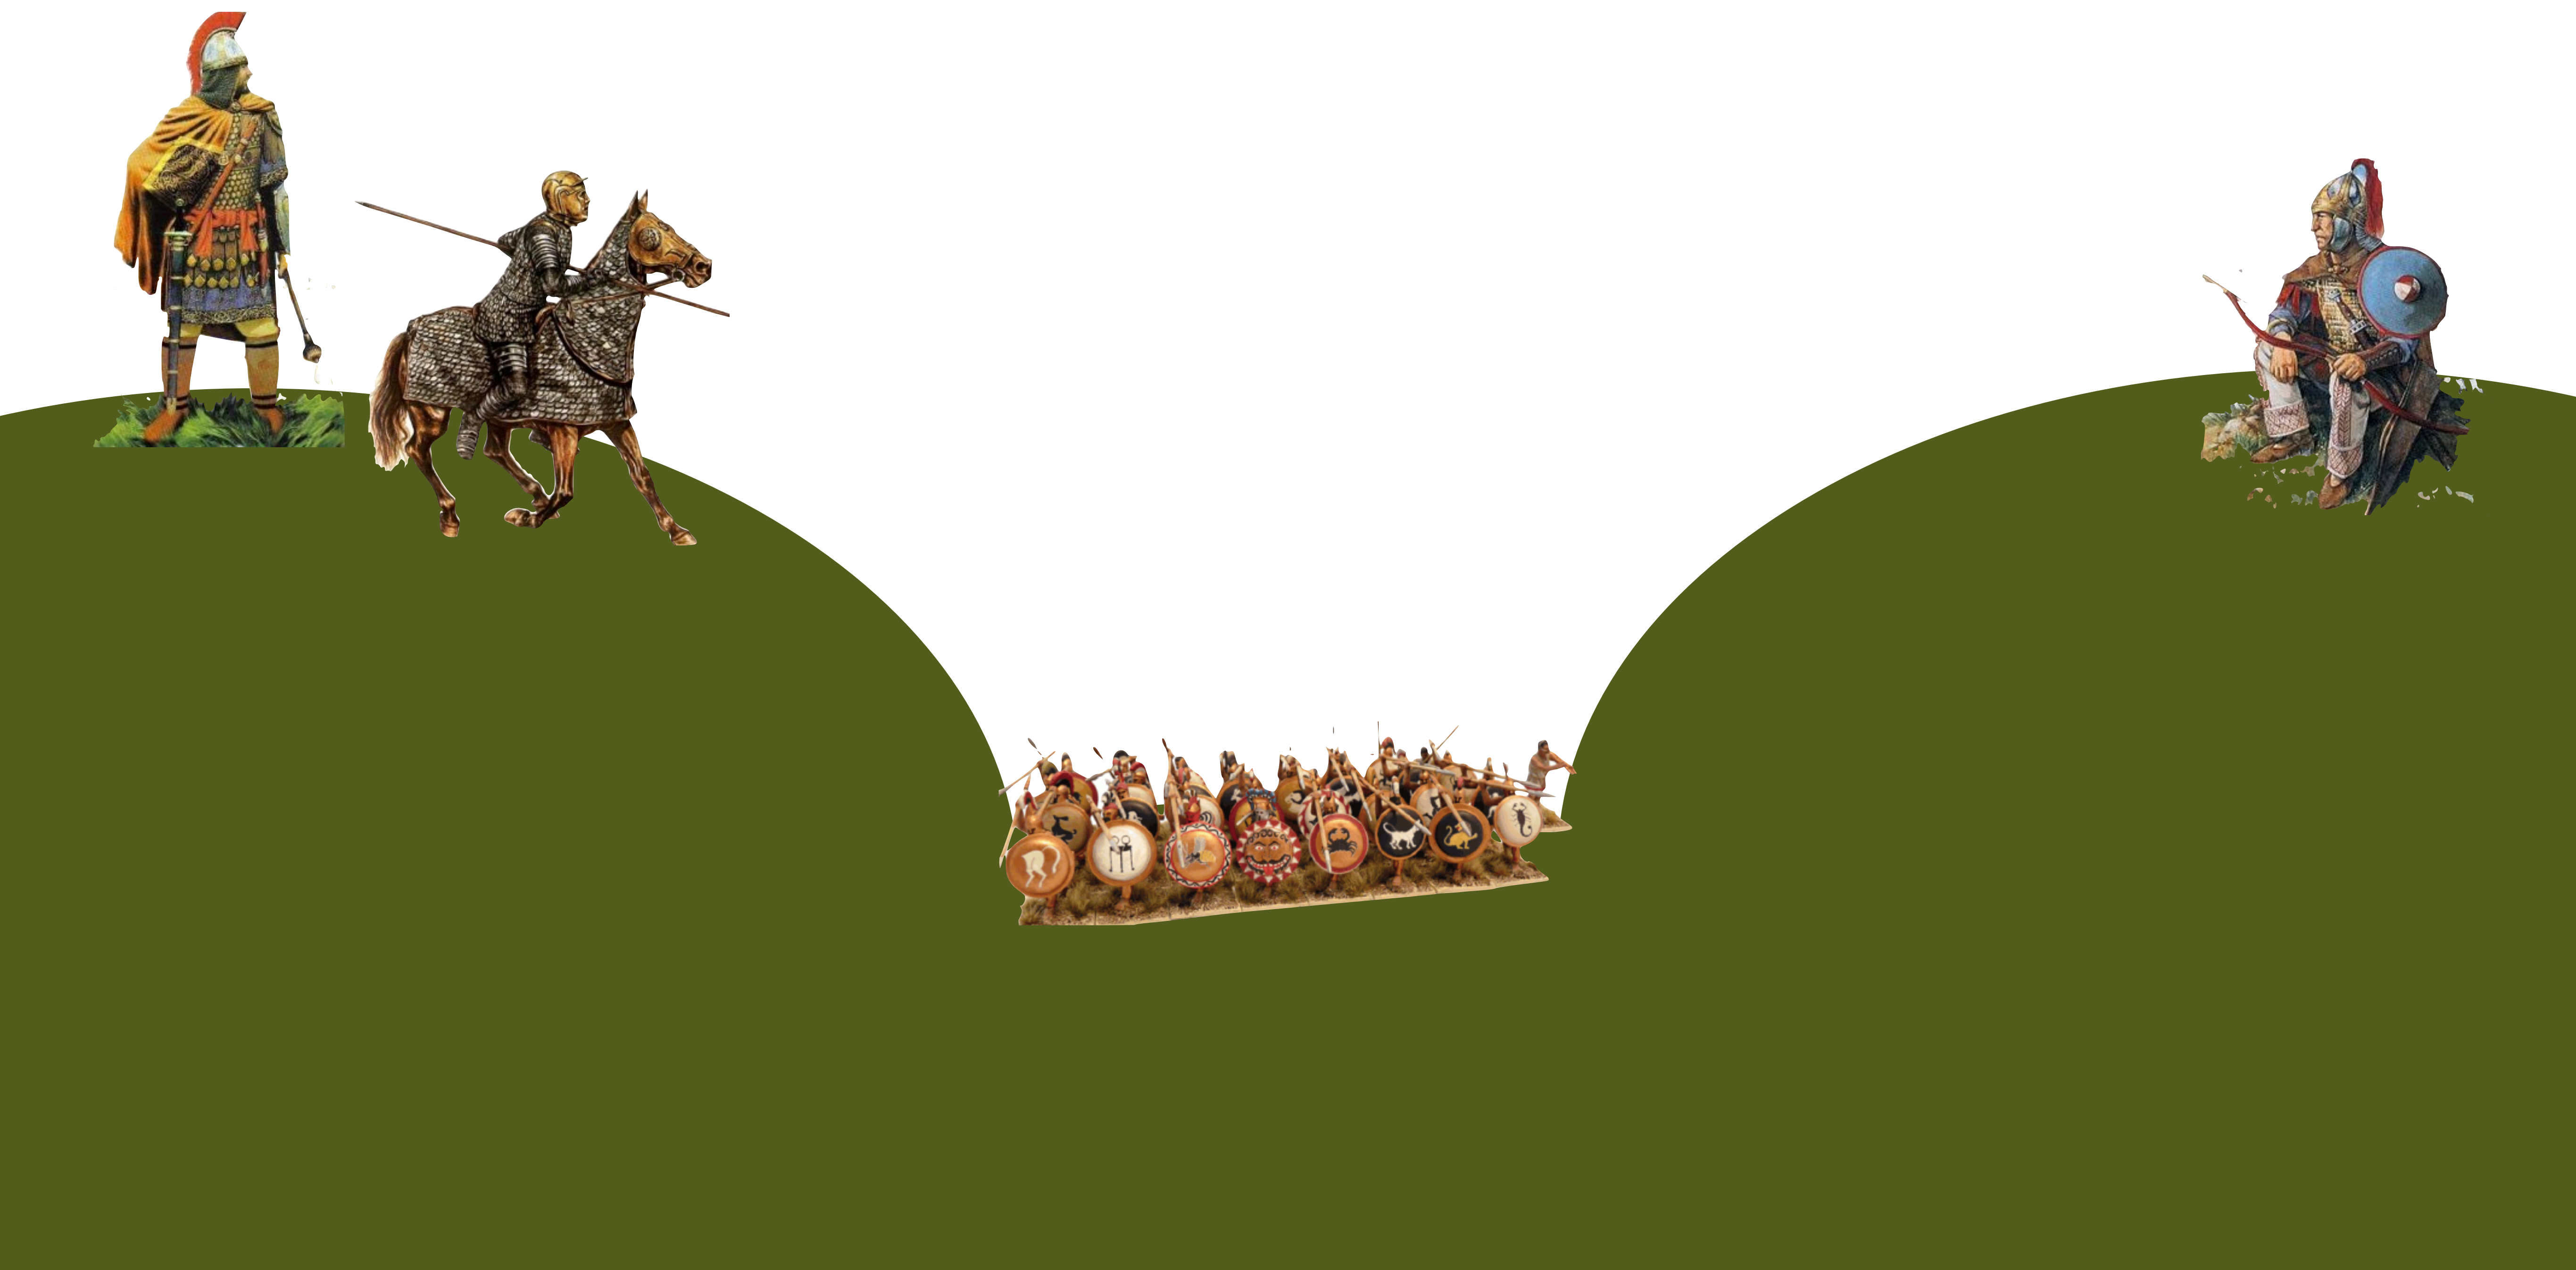


Dear General B,
The attack will be at dawn tomorrow.
Please confirm.
with love, General A.

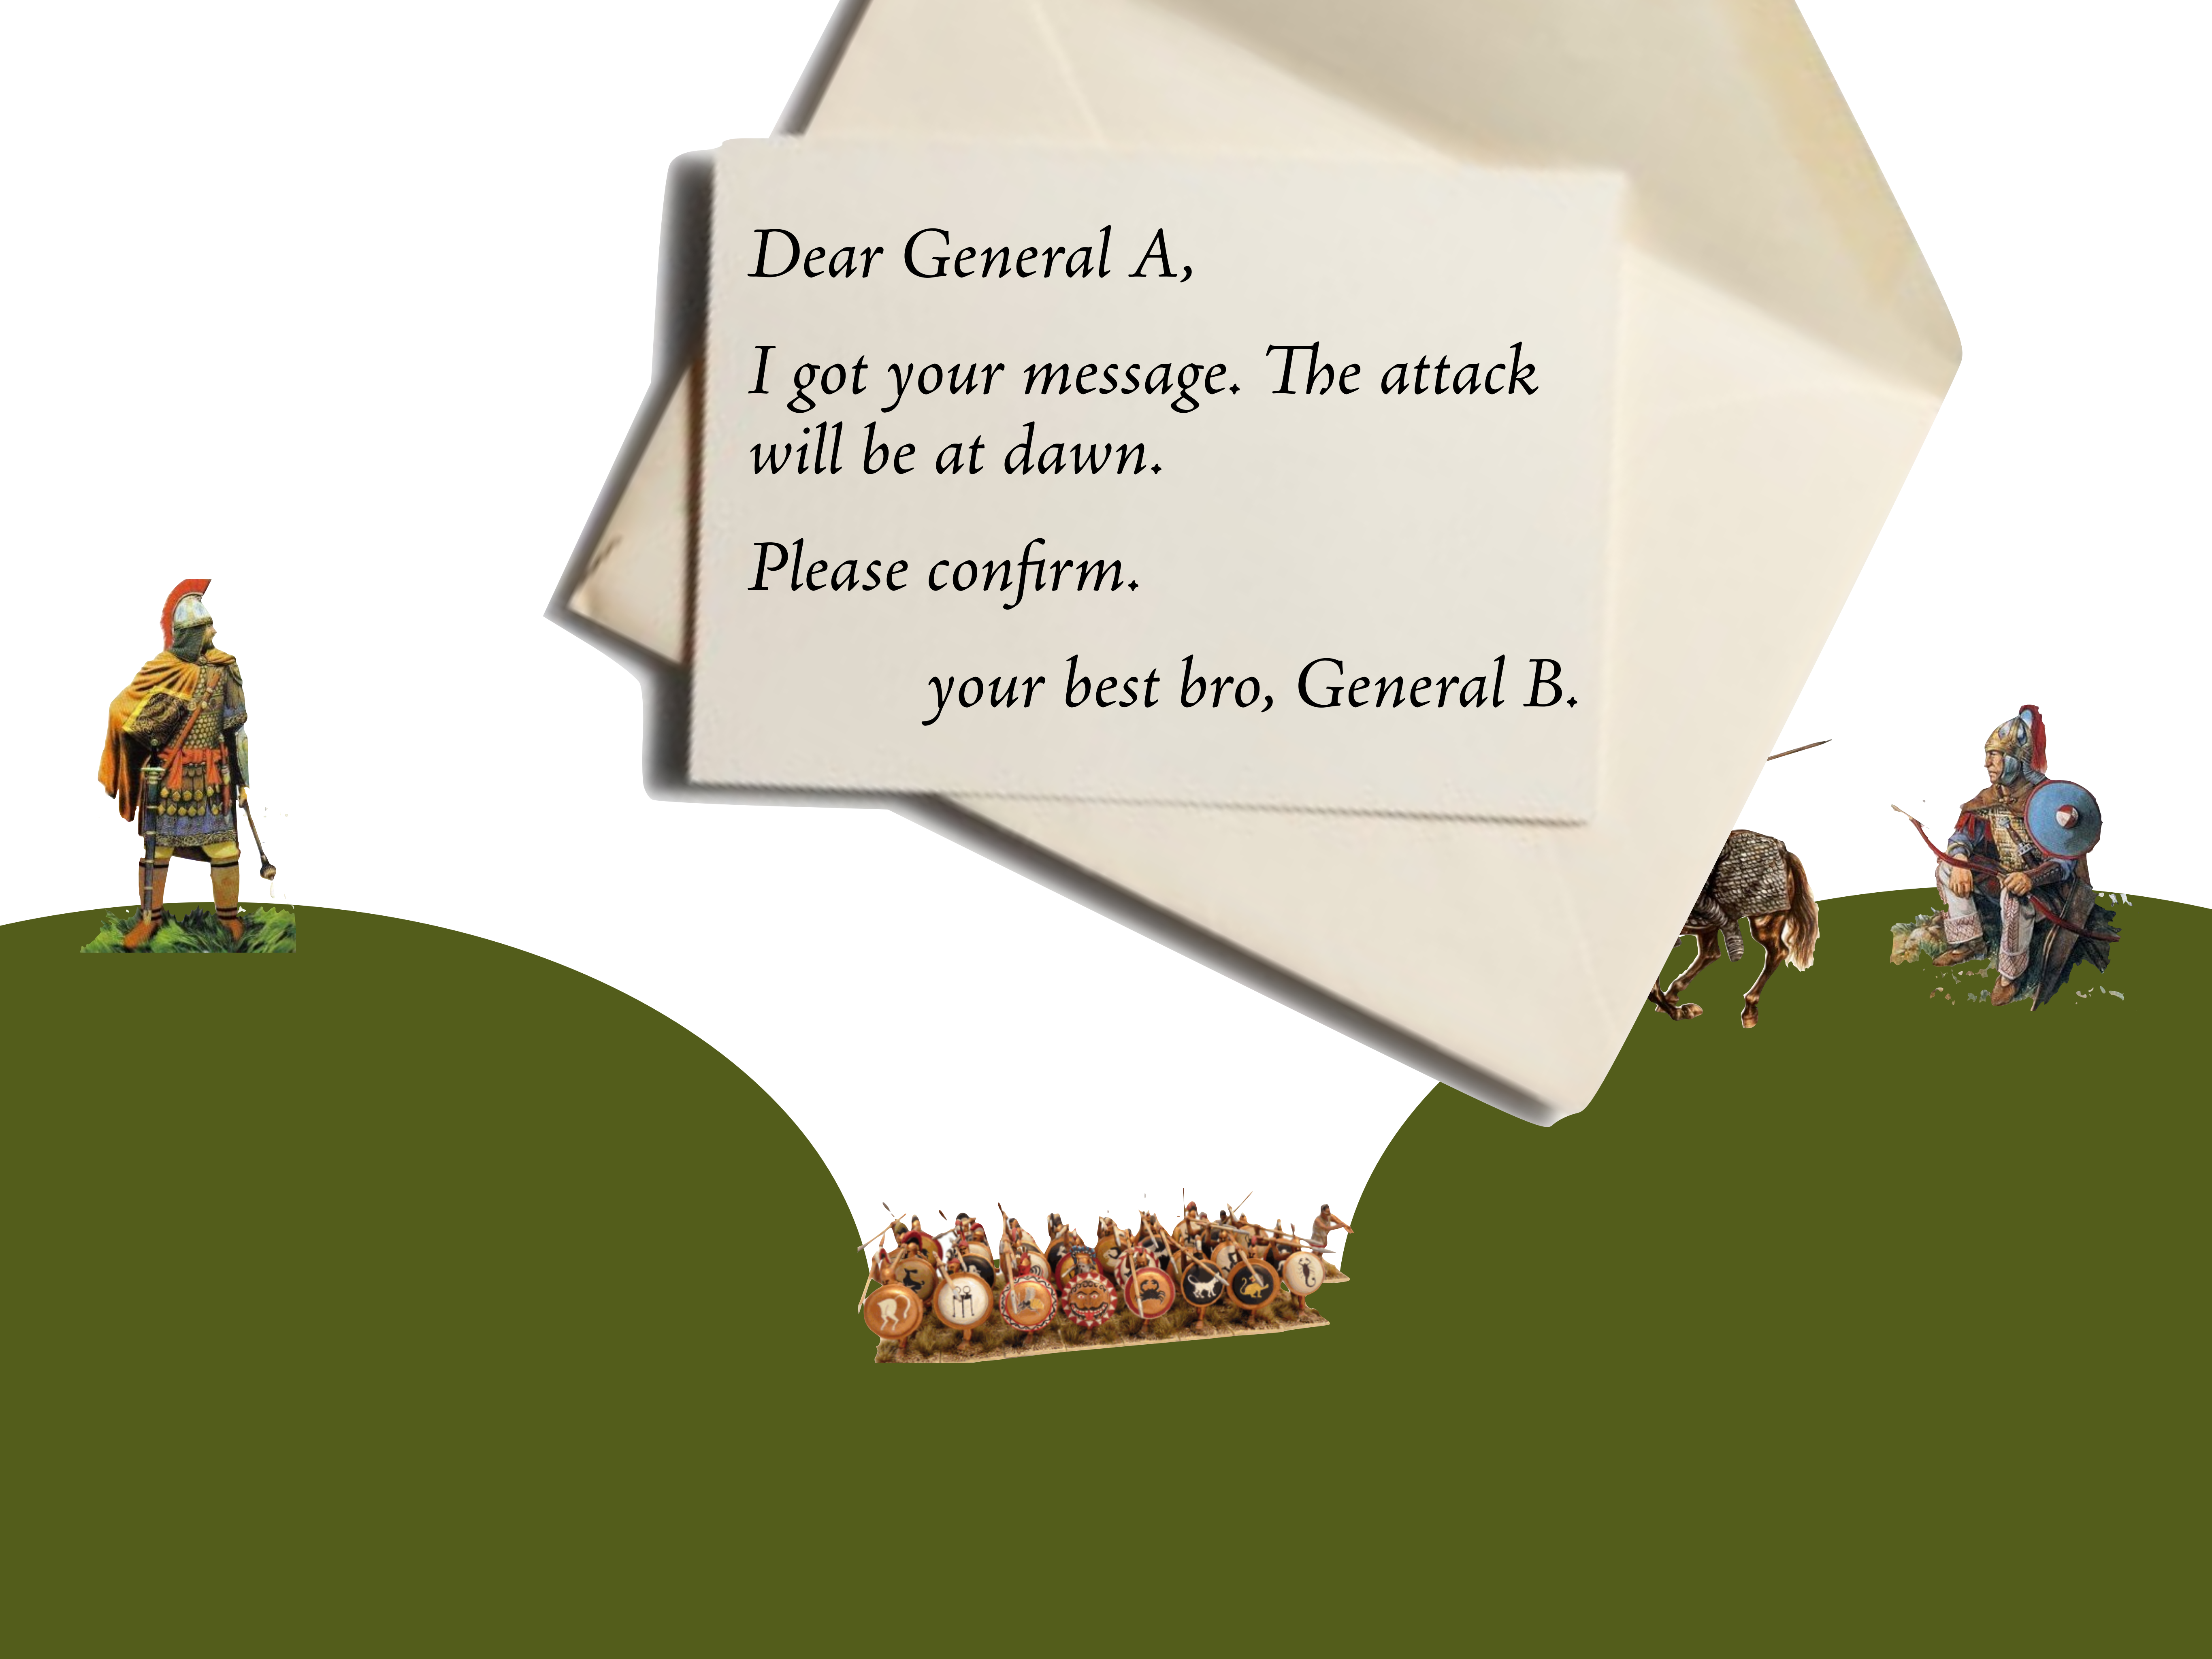
Dear General A,
I got your message. The attack will be at dawn.
Please confirm.
your best bro, General B.


Dear General B,
Got it. I love the smell of battle in the morning.
Please confirm.
bros 4 life, General A.


Dear General A,
Roger. Lock and load.
Please confirm.
bro grabs, General B.

THEOREM

Given reasonable assumptions about the generals' utility functions and epistemic standards, they will never achieve common knowledge or common belief. (Akkoyunlu et al., 1975; Gray, 1978; Halpern and Moses, 1990; Yemini and Cohen, 1979)

A (PRETTY CLEAR) COROLLARY

They won't achieve common acceptance, either.

Dear General B,

I've been reading some theoretical computer science papers, and it turns out that this is never going to work.

Anyway, my men have come down with cholera. Do you know the cure?

kisses, General A


Dear General A,
Shame about the attack.
Wash your hands and don't eat so close to the latrines.
💙💙💙, General B.


Last Will
and
Testament

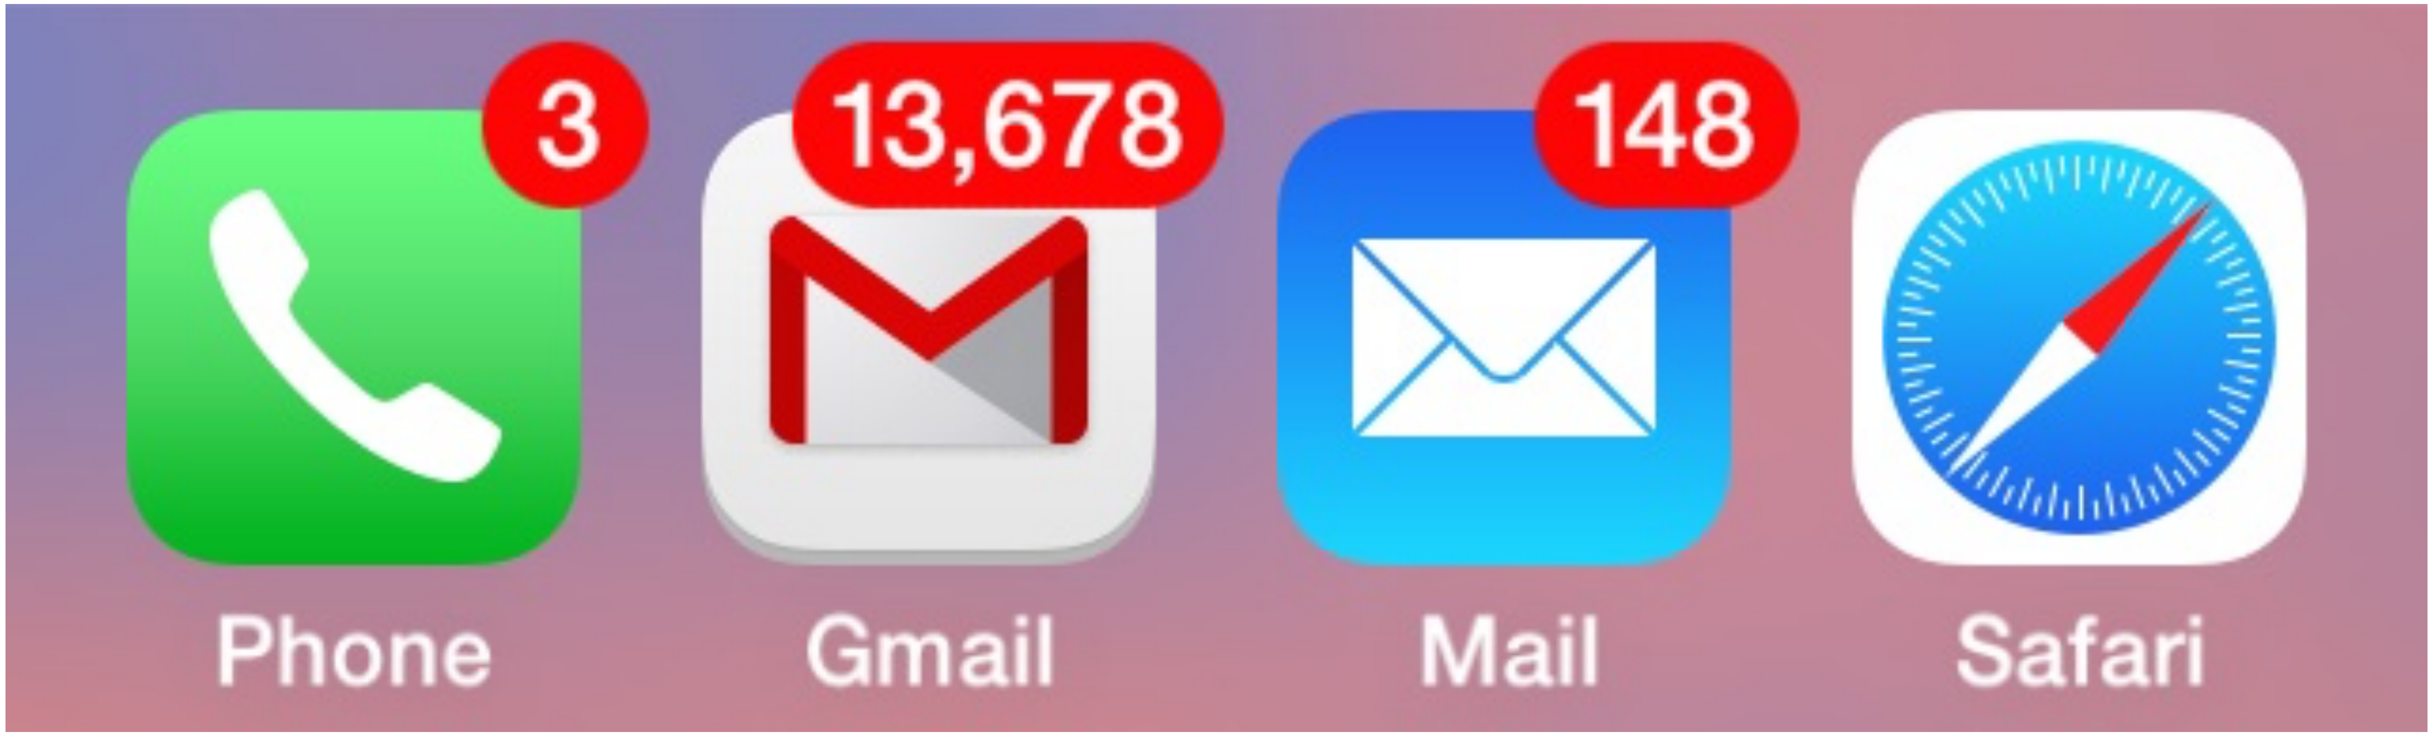

(Rubenstein 1989; Binmore 1998)

CONCLUSIONS

Successful communication doesn't require changing the context, if the context is built out of common (or even shared) attitudes.

Performing a speech act doesn't require *intending* or *proposing* to update the context, either.

CONCLUSIONS

Context change can result from communication only in certain special circumstances.

Which circumstances?

When the speaker and addressee are in a **shared situation** (Schiffer 1972; Clark & Marshall 1981).

SHARED SITUATIONS

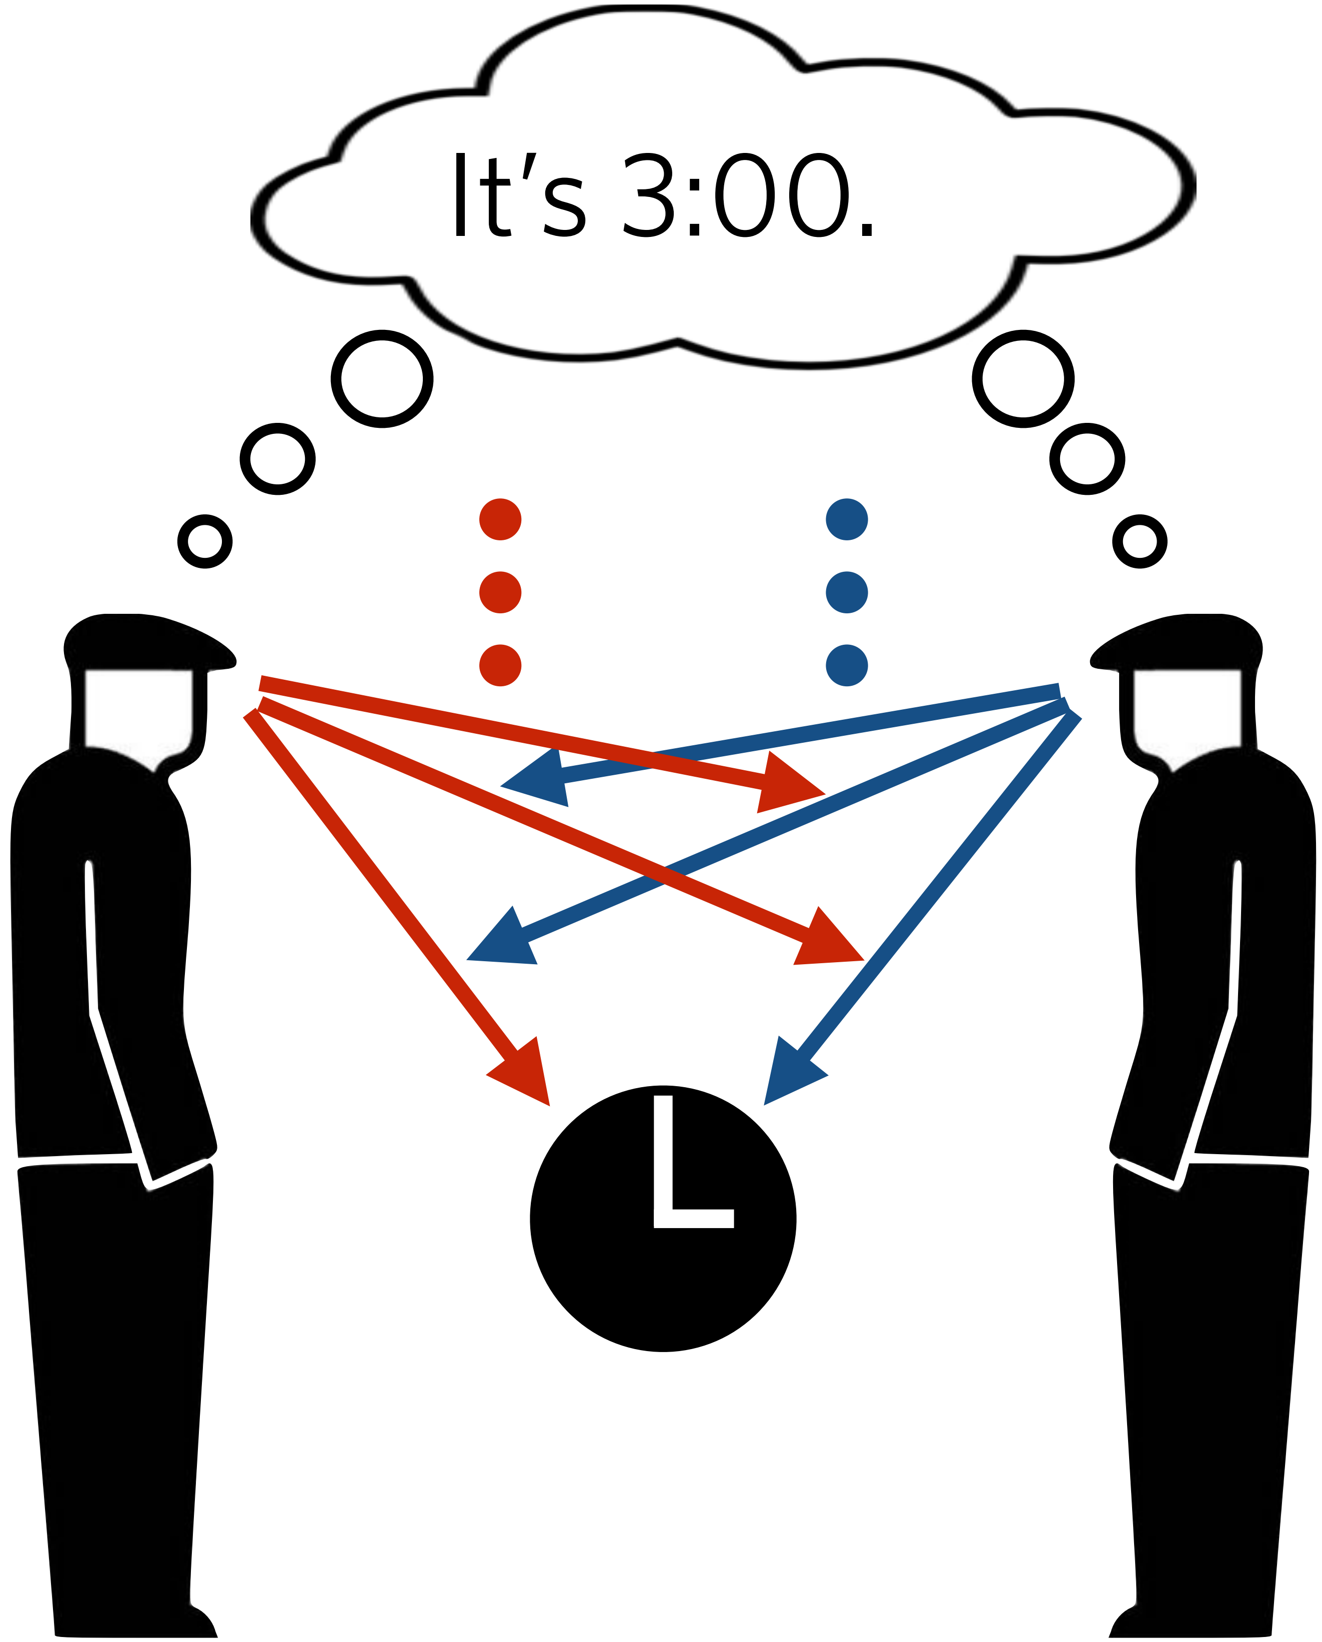

SHARED SITUATIONS

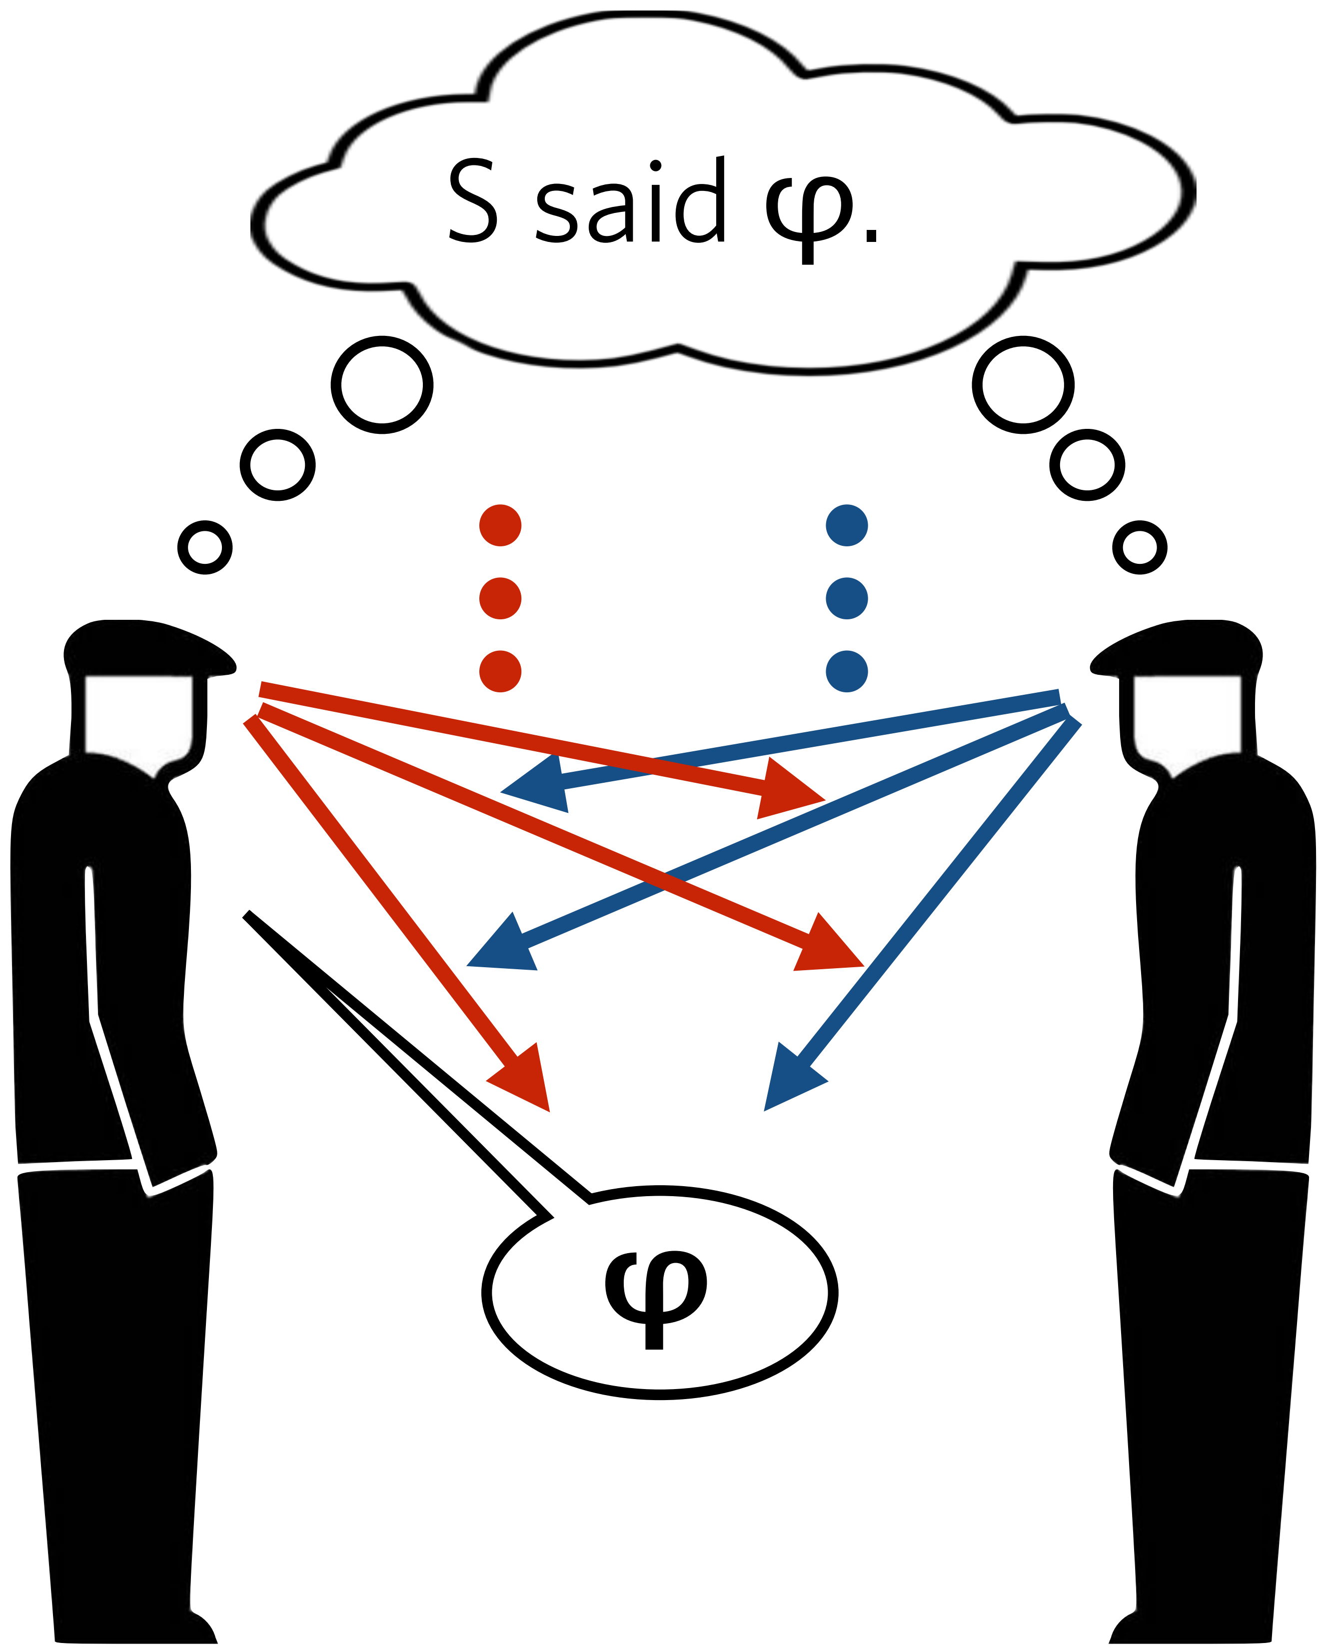

SECOND PROBLEM FOR DYNAMIC VIEWS:
BAD THEORY OF SPEECH ACTS

Solution:

Adopt a (slightly) different theory of speech acts.

My preferred option is Grice's original view.

INTENTION-BASED SEMANTICS

cf. Grice, Strawson, Schiffer, Bach & Harnish, Neale

INTENTION-BASED SEMANTICS

A PRAGMATIC VIEW

To perform a speech act is to produce an utterance with an addressee-directed communicative intention.

A METASEMANTIC VIEW

The semantic properties of expression types are to be explained in terms of the psychological states involved in using them to perform speech acts.

MEANING AND INTENDING

By doing something, x, S, **MEANT** something iff, for some audience, A, and response **R**, S did x intending

(i) A to to have a certain response **R**

(ii) A to recognise that S did x intending (1)

MEANING AND INTENDING

By doing something, x, S, **MEANT** something iff, for some audience, A, and response **R**, S did x intending

(i) A to to have a certain response **R**

(ii) A to recognise that S did x intending (1)

Different kinds of speech act are aimed at different kinds of responses

i.e., different values for **R**

MEANING AND INTENDING

By doing something, x, S, **MEANT** something iff, for some audience, A, and response **R**, S did x intending

(i) A to to have a certain response **R**

(ii) A to recognise that S did x intending (1)

To assert *p* is to communicatively intend for one's addressee to form a belief that *p*.

$$\mathbf{R} = \lambda \mathbf{M}_{\mathbf{A}} \,.\, \mathbf{B}_{\mathbf{A}} \subseteq p$$

MEANING AND INTENDING

By doing something, x, S, **MEANT** something iff, for some audience, A, and response **R**, S did x intending

(i) A to to have a certain response **R**

(ii) A to recognise that S did x intending (1)

To direct someone to ψ is to communicatively intend for them to form an intention to ψ.

$$\mathbf{R} = \lambda \mathbf{M}_{\mathbf{A}} \,.\, \mathbf{I}_{\mathbf{A}} \subseteq \{w : A\ \psi\text{s at } w\}$$

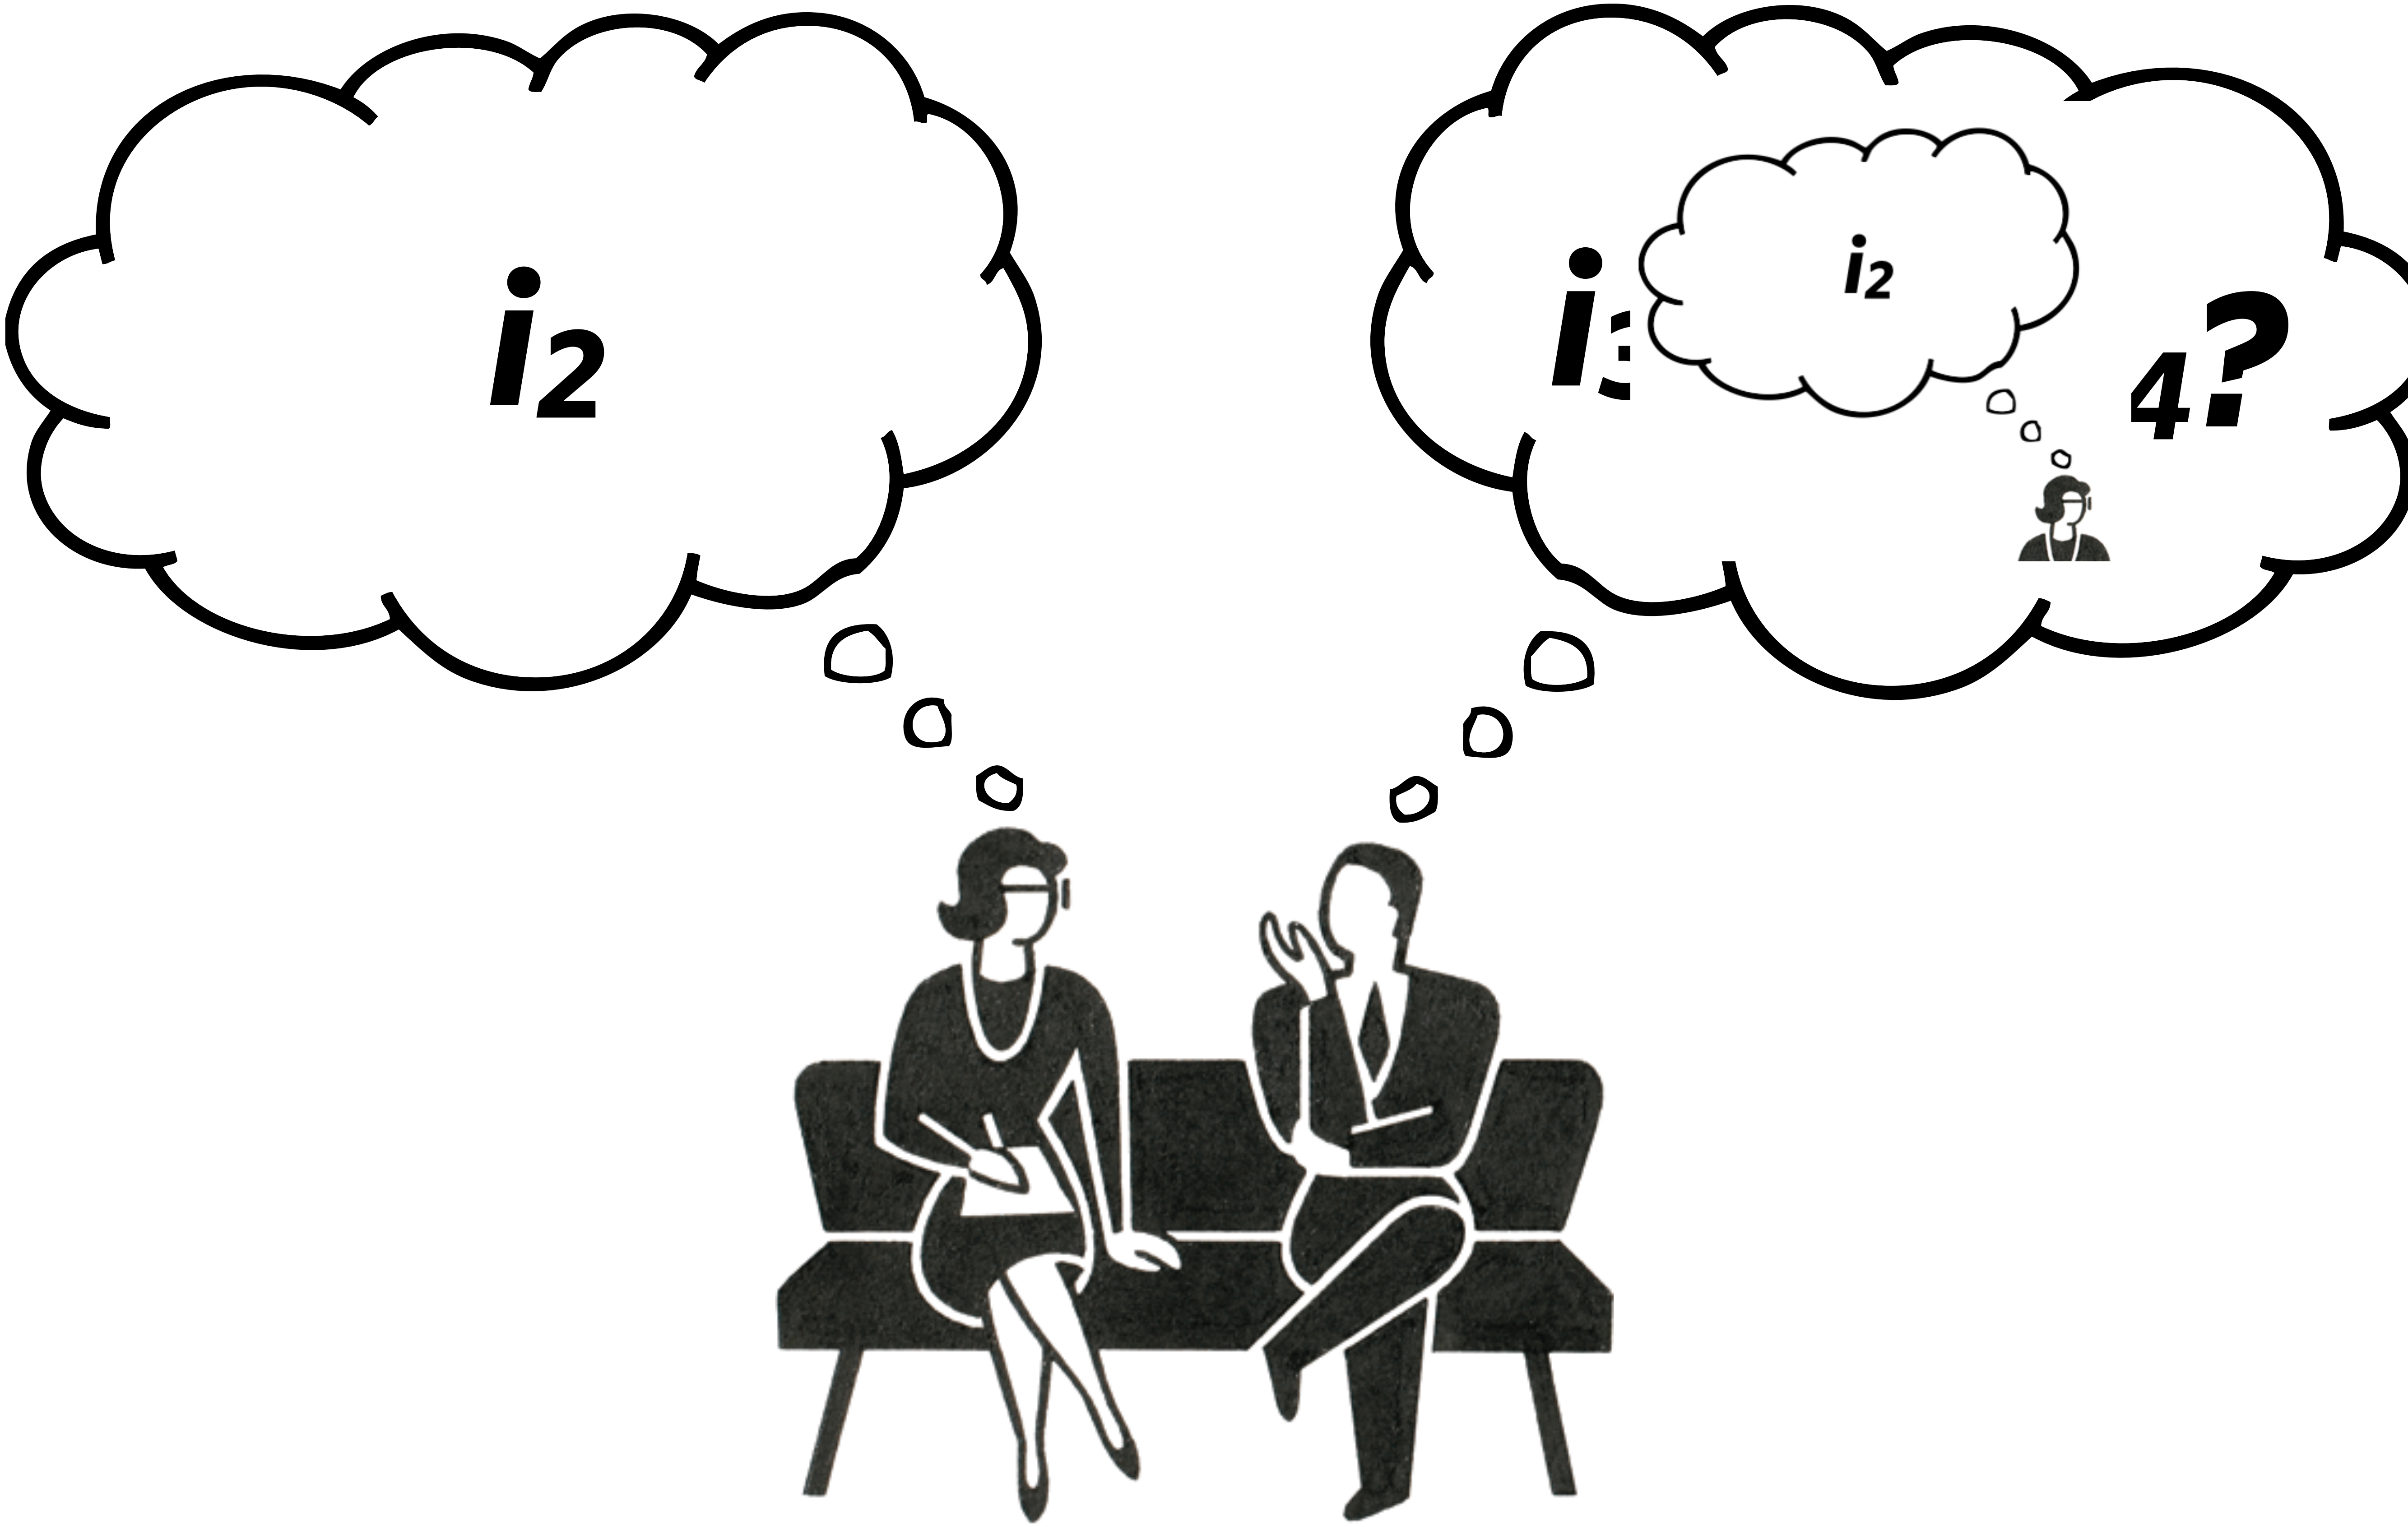

i_2
i_3
i_2
i_4?

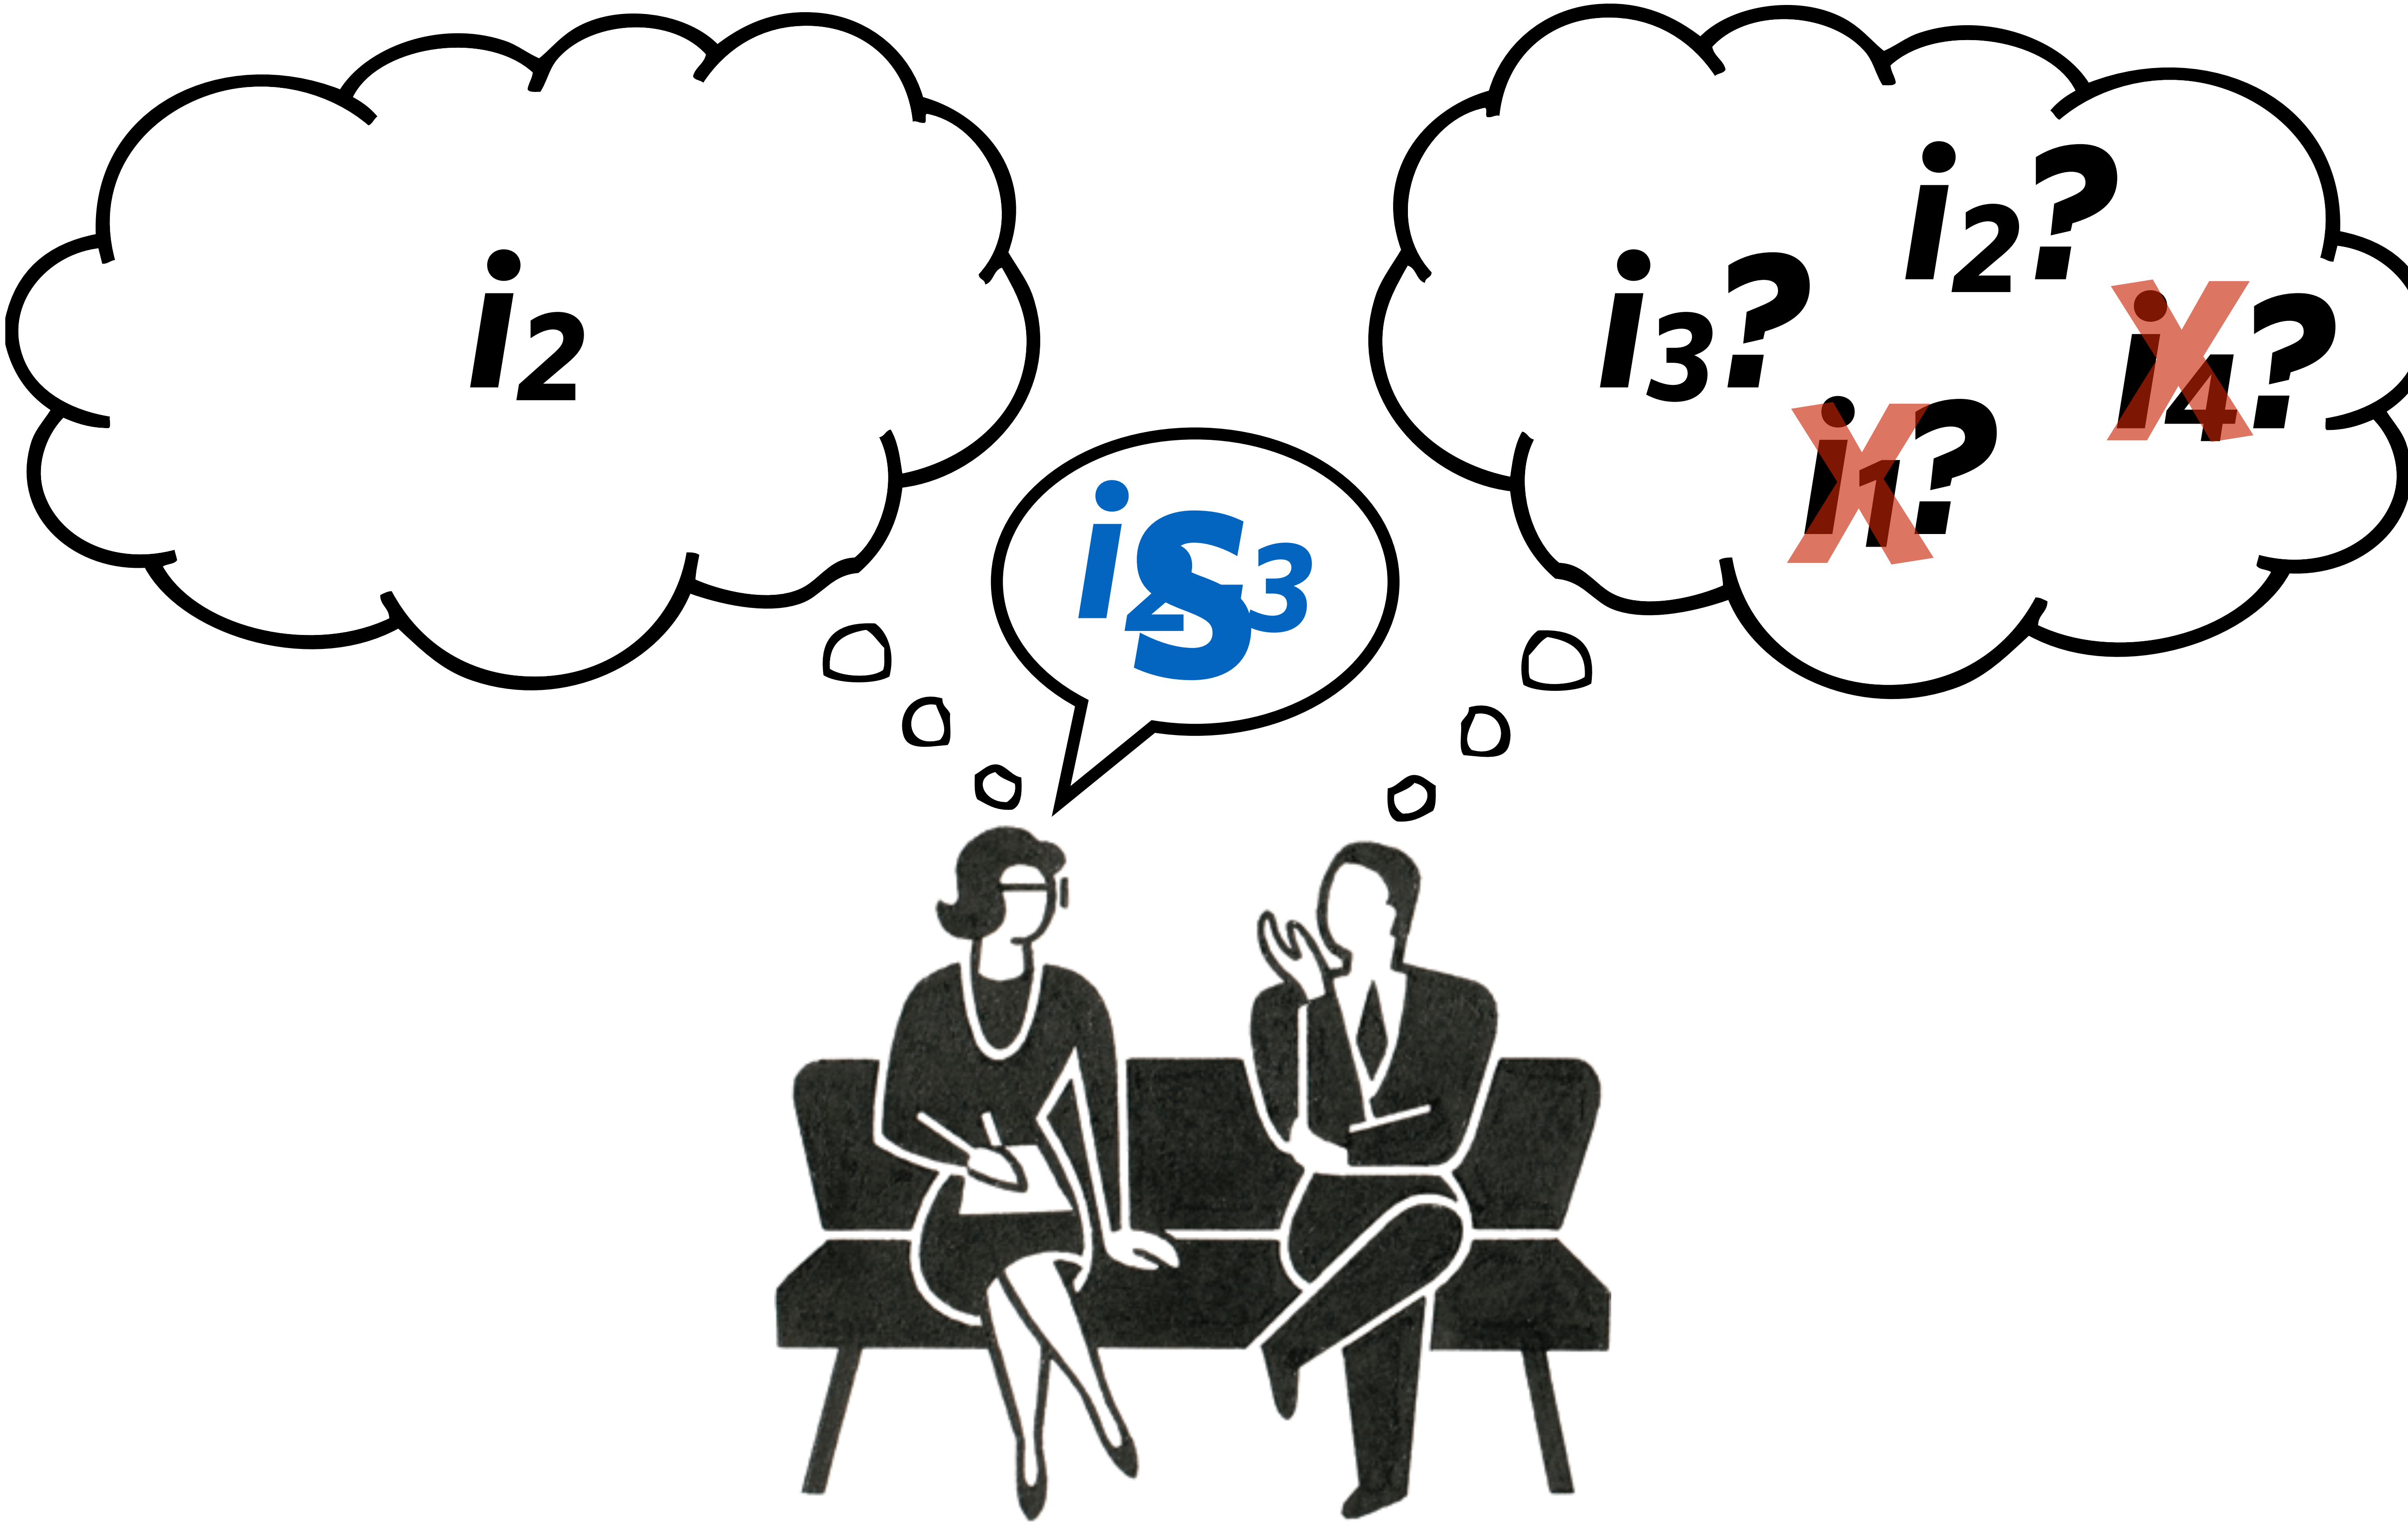

i_2
i_3?
i_2?
i_1?
i_4?

i_1?
i_2?
i_3?
i_4?
i_2

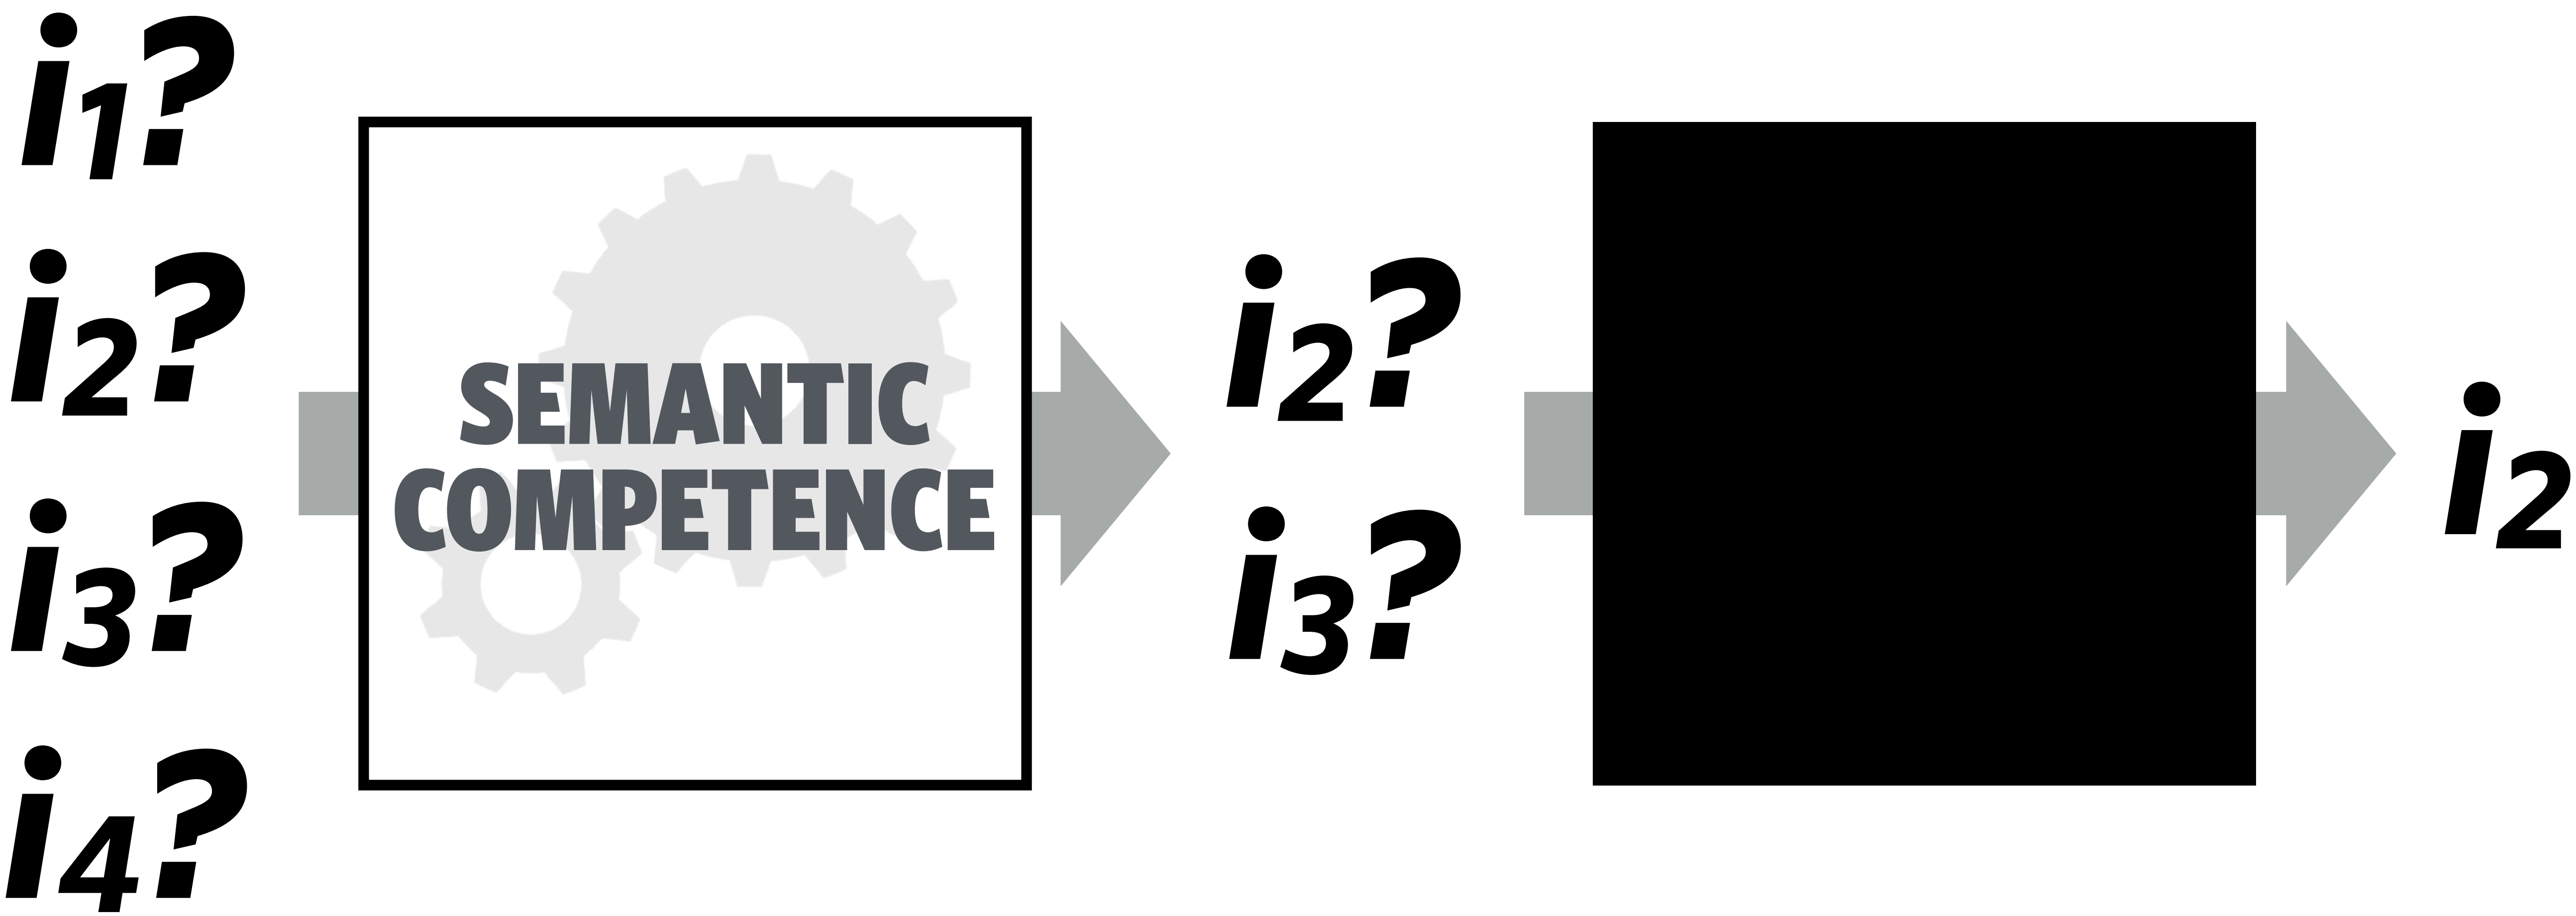
i_1?
i_2?
i_3?
i_4?
SEMANTIC
COMPETENCE
i_2?
i_3?
i_2

SEMANTICS

…is the study of a component of the mind that computes partial and defeasible evidence about what speakers intend by their utterances.

Specifically, it computes **R** values
(or at least properties of **R** values).

METASEMANTICS

A clause Φ has ⟦Φ⟧ as its semantic value for a speaker S in virtue of the fact that:

a. If S were to communicatively intend to produce ⟦Φ⟧ in an addressee, they might utter an unembedded token of Φ.

b. If S were to perceive an unembedded utterance of Φ, they would conclude that, if the speaker is being literal and direct (etc.), the speaker intends to produce ⟦Φ⟧ in their addressee.

c. (a) and (b) are true in virtue of facts about the semantic component of S's faculty of language.

METASEMANTICS

For any intuitive instance of logical consequence, $\Phi \vDash \Psi$, the fact that it strikes us as valid is explained by our sensitivity to the following fact:

> A structurally rational* agent who is in mental state $[\![\Phi]\!]$ is also in mental state $[\![\Psi]\!]$.

*A structurally rational agent is one who exemplifies Bratman-style coherence relations.

THE SEMANTICS

COGNITIVE MODELS

M_A

A's **COGNITIVE MODEL**

=a set-theoretic representation of A's beliefs and plans

M_A determines:

(in a way to be explained)

B_A

A's **BELIEF STATE**

=the set of worlds compatible with what A believes

I_A

A's **INTENTION STATE**

=the set of worlds compatible with what A intends

COGNITIVE MODELS

M_A

A's **COGNITIVE MODEL**

=a set-theoretic representation of A's beliefs and plans

...under the idealized assumption that A is structurally rational.

TWO KINDS OF IDEALIZATION
FIRST: GALILEAN

B_A I_A

Conflate mental states with necessarily equivalent contents.

(Distinction between Galilean and Minimal idealization due to Weisberg 2007)

TWO KINDS OF IDEALIZATION
SECOND: MINIMALIST

B_A I_A

Mental states are consistent and closed under entailment.

TWO KINDS OF IDEALIZATION
SECOND: MINIMALIST

$$I_A \subseteq B_A$$

TWO KINDS OF IDEALIZATION
SECOND: MINIMALIST

$$I_A \subseteq B_A$$

Doxastic Constraint on Intending

A can't intend to ψ if it is ruled out by A's beliefs that A will ψ.

TWO KINDS OF IDEALIZATION
SECOND: MINIMALIST

$$I_A \subseteq B_A$$

Strict Means-End Coherence

If:

(I) A intends to φ.
(II) A believes that ψing is necessary for φing.

then:

(III) A intends to ψ.

COGNITIVE MODELS

$$I_A \subseteq B_A$$

In other words:

We model an agent's plans as a selection function.

This function maps each belief state to a subset of itself, which is the set of worlds compatible with what the agent intends.

COGNITIVE MODELS

$$I_A \subseteq B_A$$

How to define plans as selection functions?

1. Stipulate it as a primitive
2. $M_A = \langle B_A, X \rangle$, where X is a set of intentions or preferences that imposes an ordering $<_X$ on BA. IA is $\max_{<X}(B_A)$. (Charlow 2014)
3. Construct M_A from A's belief worlds and the HYPERPLANS compatible with A's plans.

PLANS AND HYPERPLANS

A's planning state, **P_A**, is a set of **HYPERPLANS**.

A **HYPERPLAN** for A is a selection function that maps each of A's possible belief states to the intersection of that belief state with one of A's **CHOICE WORLDS**. (cf. Yalcin 2012)

A's **CHOICE WORLDS** are sets of worlds that are equivalent with respect to all of the choices that A could ever have to make.

Intuitively: each of A's **HYPERPLANS** makes every choice that A could ever have to make.

PLANS AND HYPERPLANS

$$M_A = B_A \times P_A$$

A's COGNITIVE MODEL is a set of world/hyperplan pairs.

A's BELIEF STATE is the set of all of the world coordinates.

A's PLANNING STATE is the set of all of the hyperplan coordinates.

PLANS AND HYPERPLANS

Intuitively: $\mathbf{P_A}$ is the set of A's current fully specified practical options—ways of turning possible belief states into full life plans.

$$\mathbf{I_A} = \{w : (\exists h \in \mathbf{P_A})(w \in h(\mathbf{B_A}))$$

A's **INTENTION STATE** is the set of worlds in A's **BELIEF STATE** not ruled out by all of A's current practical options.

CLAUSAL SEMANTICS

The semantic values of clauses are properties of cognitive models:

For any sentence radical φ,

$$[\![\vdash\varphi]\!]^c = \lambda \mathbf{M}_A\,.\,\mathbf{B}_A \subseteq [\![\varphi]\!]^c$$

$$[\![!\varphi]\!]^c = \lambda \mathbf{M}_A\,.\,\mathbf{I}_A \subseteq [\![\varphi]\!]^c$$

PROBLEM FOR STATIC VIEWS:
MIXED COORDINATION

(We're about to go into the bar together. I say:)

Buy us drinks and I'll find a table.

NOTE:

- Needn't have a conditional meaning.
- Can mean roughly: 'I'll find a table. Buy me a drink.'
- Can be the consequent of a conditional:
  'If your friend is tending bar, buy us drinks and I'll find a table.'

CONJUNCTION

Where Φ,Ψ are sentences that may be imperatives, declaratives, or combinations of the two:

What is ⟦ Φ and Ψ⟧?

CONJUNCTION

$$[\![\Phi \text{ and } \Psi]\!] =$$

$$\lambda M \,.\, [\![\Phi]\!](M) = 1 \text{ and } [\![\Psi]\!](M) = 1$$

CONJUNCTION

⟦Buy us drinks and I'll find a table⟧^{c} =

λM . M intends to buy drinks and M believes that speaker$_c$ will find a table

(The property of being a mind that intends to buy drinks and believes that speaker$_c$ find a table.)

PROBLEM FOR STATIC VIEWS:
MIXED COORDINATION

(We're at a book store. Each of us has three books, but we only have enough money for five, total:)

Put back Naked Lunch or I'll put back Waverley.

(Starr ms)

NOTE:

- Needn't have a conditional meaning.
- Can be the consequent of a conditional: 'If we only have $5, put back Naked Lunch or I'll put back Waverly.'

DISJUNCTION

Where Φ,Ψ are sentences that may be imperatives, declaratives, or combinations of the two:

What is $[\![\Phi \text{ or } \Psi]\!]$?

DISJUNCTION

WEAK DISJUNCTION

$$[\![\Phi \text{ or } \Psi]\!] = \lambda M\,.\,(\exists M^1, M^2)\begin{pmatrix} [\![\Phi]\!](M^1) = 1 \text{ and} \\ [\![\Psi]\!](M^2) = 1 \text{ and} \\ M^1 \cup M^2 := M \end{pmatrix}$$

DISJUNCTION

WEAK DISJUNCTION

$[\![$Put back Naked Lunch or I'll put back Waverly$]\!]^c$ =

λM . M is in a state of either:

(a) intending to put back Naked Lunch; or

(b) believing that speakers will put back Waverly; or

(c) indecision between options (a) and (b), but commitment to at least one.

DISJUNCTION

A problem with **WEAK DISJUNCTION**:

(1) A: I'll put back Naked Lunch.

(2) B: Put back Naked Lunch *or* I'll put back Waverly.

If A is being sincere with (1), then B's response is redundant. So why doesn't it *sound* redundant?

DISJUNCTION

A problem with **WEAK DISJUNCTION**:

(1) A: Class is in room 505.

(2) B: It’s in 505 *or* it’s in 506.

If A is being sincere with (1), then B’s response is redundant. So why doesn’t it *sound* redundant?

DISJUNCTION

WEAK DISJUNCTION

$[\![$Put back Naked Lunch or I'll put back Waverly$]\!]^c$ =

λM . M is in a state of either:

(a) intending to put back Naked Lunch; or

(b) believing that speakers will put back Waverly; or

(c) indecision between options (a) and (b), but commitment to at least one.

DISJUNCTION

WEAK DISJUNCTION

⟦Put back Naked Lunch or I'll put back Waverly⟧^{c} =

λM . M is in a state of either:

(a) ~~intending to put back Naked Lunch; or~~

(b) ~~believing that speakers will put back Waverly; or~~

(c) indecision between options (a) and (b), but commitment to at least one.

DISJUNCTION

STRONG DISJUNCTION

$$[\![\Phi \text{ or } \Psi]\!] =$$

$$\lambda M\,.\,(\exists M^1 : I_{M1} \neq \varnothing)\ (\exists M^2 : I_{M2} \neq \varnothing) : \left(\begin{array}{l} [\![\Phi]\!](M^1) = 1 \text{ and} \\ [\![\Psi]\!](M^2) = 1 \text{ and} \\ M^1 \cup M^2 := M \end{array} \right)$$

Intuitively: In uttering a disjunction, I intend for you to take both alternatives seriously, at least initially and for the purposes of practical reasoning.

DISJUNCTION

STRONG DISJUNCTION

$[\![$Put back Naked Lunch or I'll put back Waverly$]\!]^c$ =

λM . M is in a state of indecision between intending to put back Naked Lunch and believing that speaker$_c$ will put back Waverly, but commitment to at least one of these options.

CONDITIONAL

Where Φ is a declarative and Ψ is a declarative, imperative, or combination of the two:

What is $[\![\text{if } \Phi \text{ then } \Psi]\!]$?

CONDITIONAL

MAXIMAL SUBMODEL

A maximal **Φ**-supporting submodel M^{Φ} of M meets the following conditions:

(i) $M^{\Phi} \subseteq M$;

(ii) $[\![\Phi]\!](M^{\Phi}) = 1$;

(iii) There is no M^{*} such that:

- $M^{\Phi} \subseteq M^{*} \subseteq M$; and
- $[\![\Phi]\!](M^{*}) = 1$

CONDITIONAL

$$[\![\text{if } \Phi \text{ then } \Psi]\!] = \lambda M \,.\, (\forall M^{\Phi})[\![\Psi]\!](M^{\Phi})$$

M satisfies $[\![\text{if } \Phi \text{ then } \Psi]\!]$ iff every maximal Φ-satisfying submodel of M satisfies Ψ.

Intuitively:

In uttering 'if Φ then Ψ', I intend you to enter a state of mind such that, if you were to also form the belief $[\![\Phi]\!]$, it would be irrational for you not to also enter the state $[\![\Psi]\!]$.

CONDITIONAL

⟦If Quinn is bartending, buy the first round⟧ =

λM_A . If A were to be in mental state M_A and believe that Quinn is bartending, A would intend to buy the first round.

CONSEQUENCE (QUICKLY)

CONSEQUENCE (QUICKLY)

$\{\Phi_1 \ldots \Phi_n\} \vDash \Psi$ iff:

$(\forall M)$ if $[\![\Phi_1]\!](M)=1, \ldots, [\![\Phi_n]\!](M)=1$, then $[\![\Psi]\!](M)=1$

Ψ follows from $\{\Phi_1 \ldots \Phi_n\}$ iff every cognitive model that satisfies all of the premises also satisfied the conclusion.

CONSEQUENCE (QUICKLY)

Buy me a drink.

You won't buy me a drink unless you go to the bar.

$\models$ So, go to the bar!

Attack if the weather is good.

The weather is good.

$\models$ So, attack!

CONSEQUENCE (QUICKLY)

ROSS'S PARADOX

Post the letter

$\nvDash$ Post the letter or burn the letter.

*Note: this inference is blocked only if we adopt **STRONG DISJUNCTION**.

CONSEQUENCE (QUICKLY)

FREE CHOICE PERMISSION

Have tea or coffee.
$\models$ Have tea.

Sounds valid only if the conclusion is read as a weak use of the imperative: a permission, acquiescence, invitation, or instruction.

Here's clause for weak imperatives that would validate this inference:

$$[\![\mathbf{i}\varphi]\!]^c = \lambda \mathbf{M}_A \,.\, \mathbf{I}_A \cap [\![\varphi]\!]^c \neq \varnothing$$

Intuitively: in saying 'have tea', my aim was for you to make having tea compatible with your plans (at least tentatively and for the purposes of practical reasoning).

THANKS