Algorithmic Game Theory and Applications

Single-Parameter Domains

Domains

Unrestricted Domain

Domains

Unrestricted Domain

Gibbard-Satterthwaite 73-75

Domains

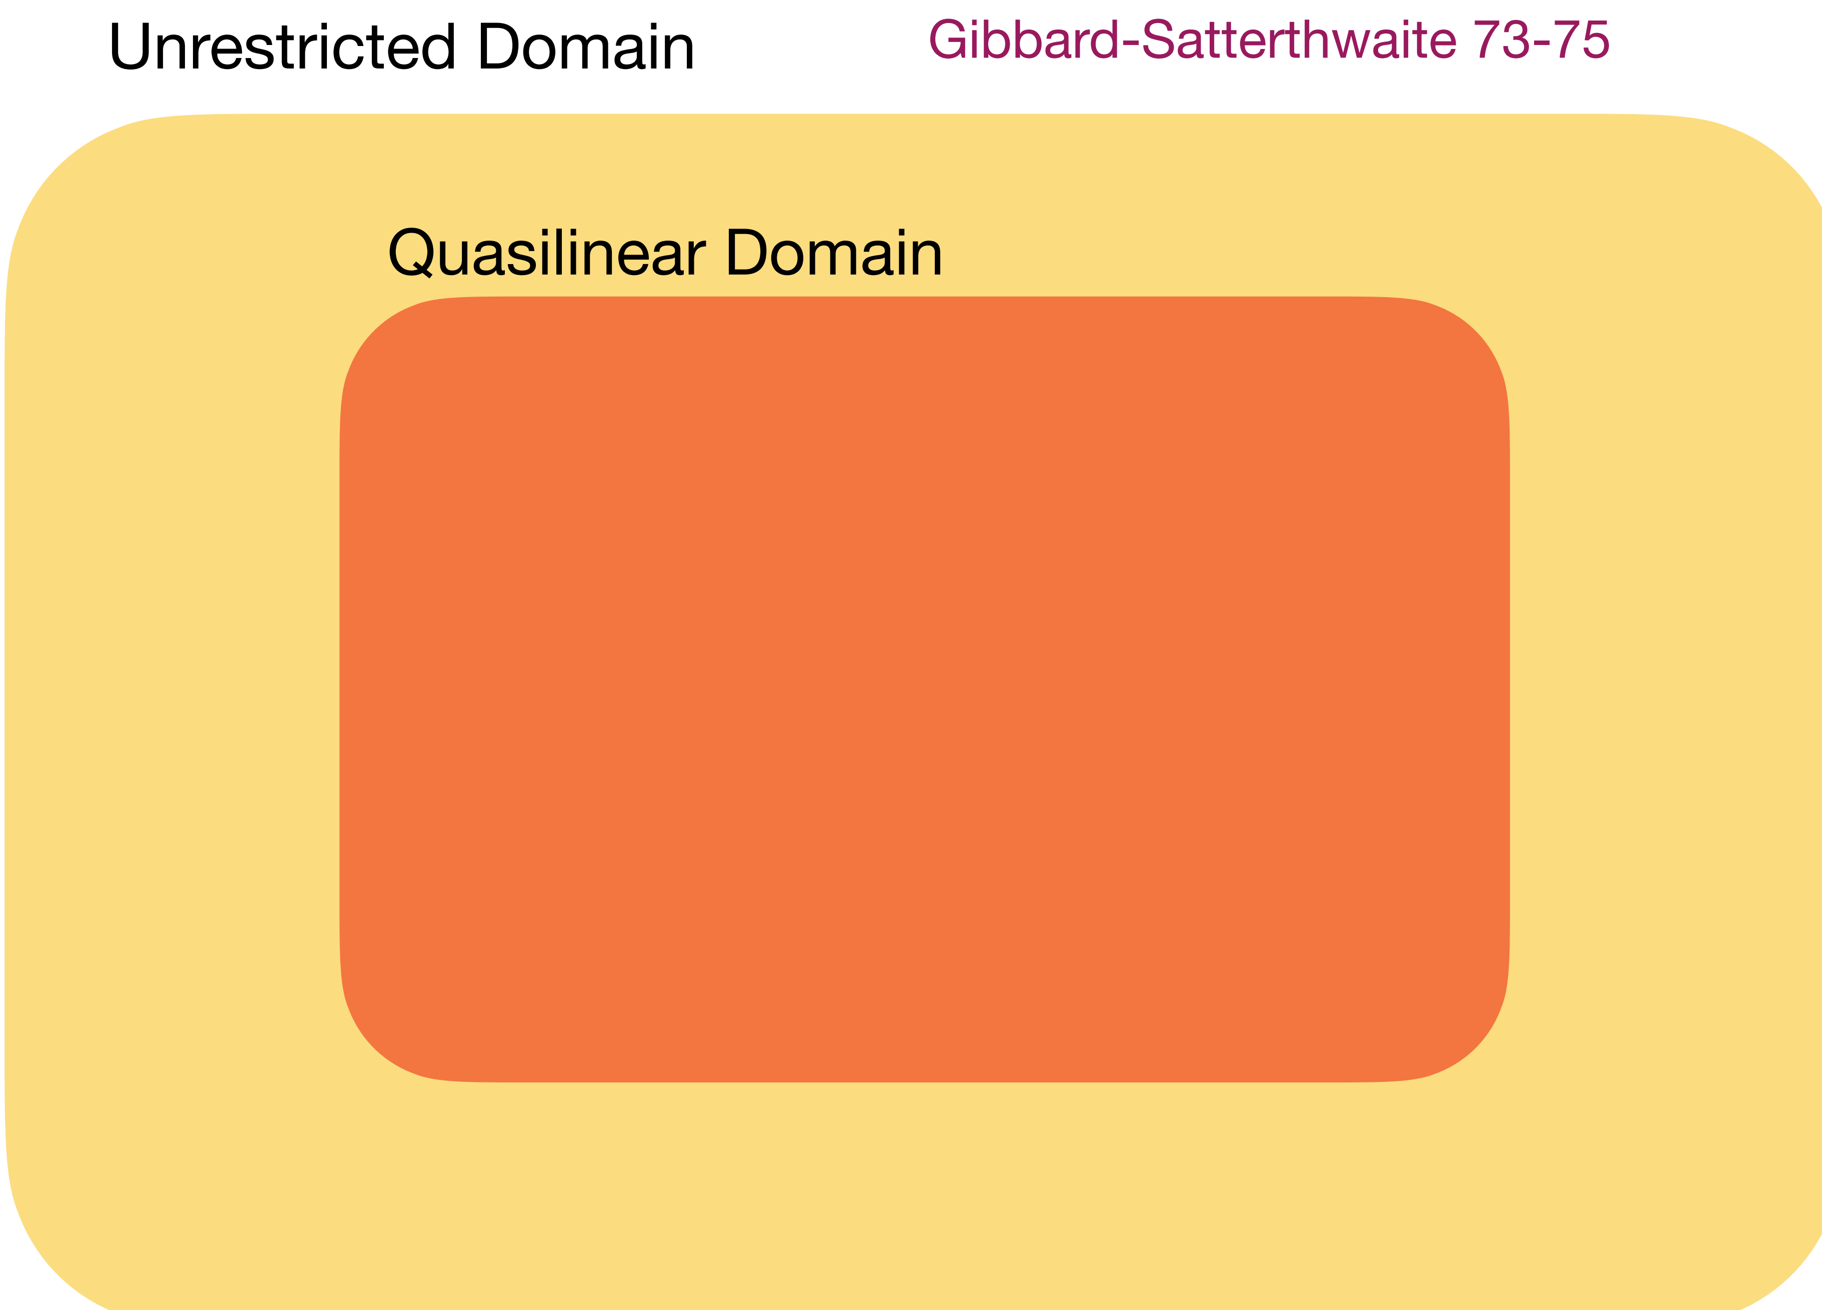
Unrestricted Domain
Gibbard-Satterthwaite 73-75
Quasilinear Domain

Domains

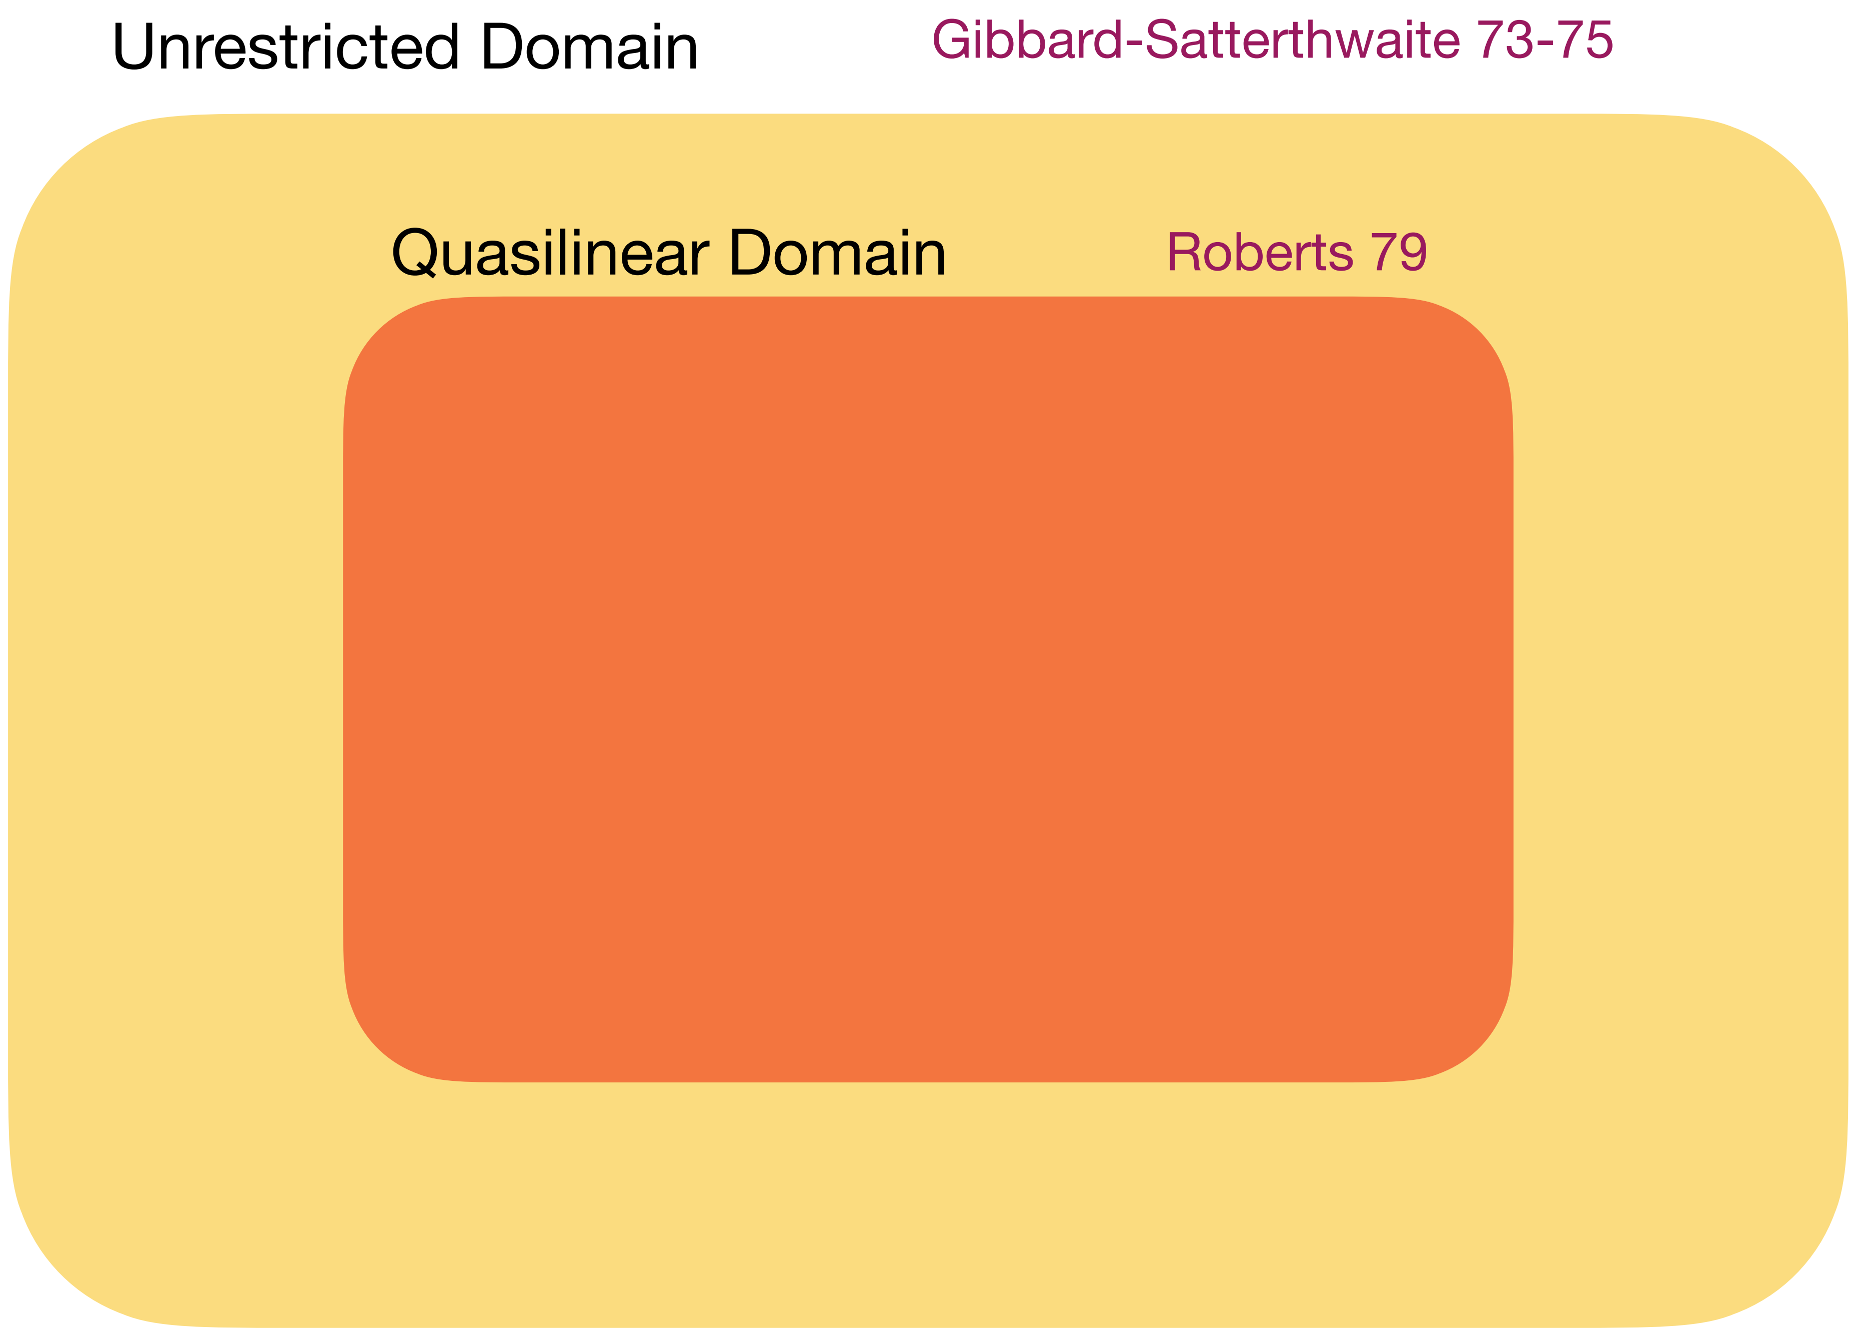
Unrestricted Domain
Gibbard-Satterthwaite 73-75
Quasilinear Domain
Roberts 79

Domains

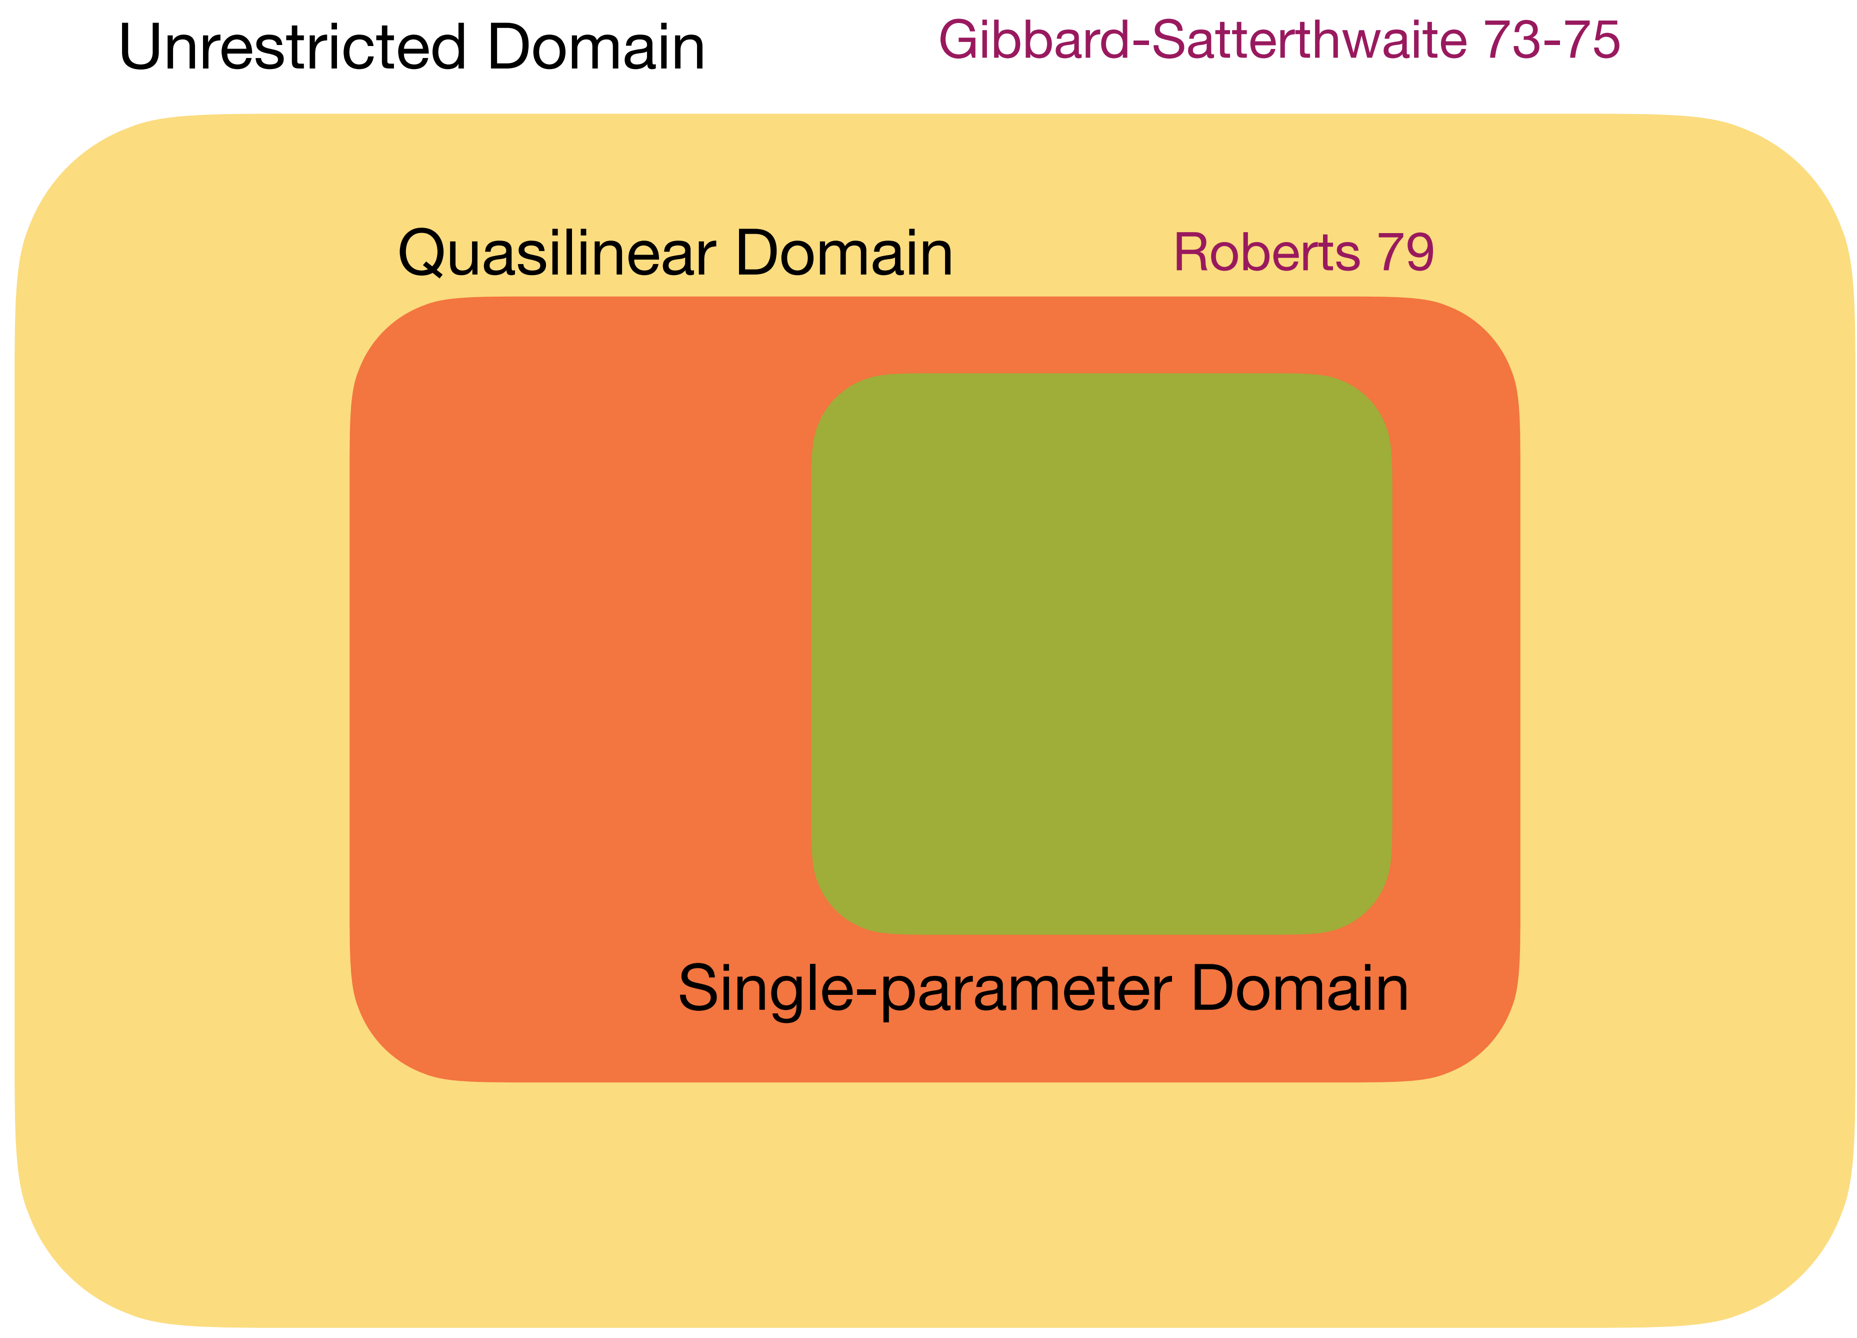
Unrestricted Domain
Gibbard-Satterthwaite 73-75
Quasilinear Domain
Roberts 79
Single-parameter Domain

Single-item Auctions

There are n bidders from a set $N = \{1,\ldots,n\}$.

There is one item for sale.

Every bidder has a value v_i for the item - this is the bidder's willingness to buy it.

Each bidder chooses a bid $b_i = \beta(v_i)$ according to some function β.

The allocation function $f : B^n \rightarrow \{0,1\}^n$ decides who wins given the bids.

The payment function $p : B^n \rightarrow \mathbb{R}^n$ decides how much each bidder will pay.

Single-item Auctions

There are n bidders from a set $N = \{1, \ldots, n\}$.

There is one item for sale.

Every bidder has a value v_i for the item - this is the bidder's willingness to buy it.

Each bidder chooses a bid $b_i = \beta(v_i)$ according to some function β.

The allocation function $f : B^n \to \{0,1\}^n$ decides who wins given the bids.

The payment function $p : B^n \to \mathbb{R}^n$ decides how much each bidder will pay.

Other examples

Other examples

Multi-unit auctions: all the items are identical and the values depend on the number of items only, i.e., $v_i(|S|)$.

Other examples

Multi-unit auctions: all the items are identical and the values depend on the number of items only, i.e., $v_i(|S|)$.

the bidder only specifies the value for one item v_i, and the value for a set S is $|S| \cdot v_i$.

Other examples

Multi-unit auctions: all the items are identical and the values depend on the number of items only, i.e., $v_i(|S|)$.

> the bidder only specifies the value for one item v_i, and the value for a set S is $|S| \cdot v_i$.

Bilateral trade: one item for sale, the seller has a value of v_s for the item and the buyer has a value of v_b. The possible outcomes are {trade, no-trade} and appropriate payments can be chosen.

Other examples

Multi-unit auctions: all the items are identical and the values depend on the number of items only, i.e., $v_i(|S|)$.

the bidder only specifies the value for one item v_i, and the value for a set S is $|S| \cdot v_i$.

Bilateral trade: one item for sale, the seller has a value of v_s for the item and the buyer has a value of v_b. The possible outcomes are {trade, no-trade} and appropriate payments can be chosen.

Public project: A public project with cost C is to be done, which is valued by each citizen at v_i. The government wants to implement the project if $\sum_i v_i > C$.

Single-parameter domains

Single-parameter domains

Informally: A domain is single-parameter if the value of each agent for the possible outcomes can be captured (encoded) by a single value v_i.

Single-parameter domains

Informally: A domain is single-parameter if the value of each agent for the possible outcomes can be captured (encoded) by a single value v_i.

Formally: There is a set of "winning outcomes" W_i for agent i, and $v_i(a) \in [t^0, t^1]$ if $a \in W_i$ and $v_i(a) = 0$, otherwise.

Single-parameter domains

Informally: A domain is single-parameter if the value of each agent for the possible outcomes can be captured (encoded) by a single value v_i.

Formally: There is a set of "winning outcomes" W_i for agent i, and $v_i(a) \in [t^0, t^1]$ if $a \in W_i$ and $v_i(a) = 0$, otherwise.

It may make sense to think of single-item auctions, keeping in mind that the results that we will present next are much more general.

Single-parameter domains

Single-parameter domains

Task: We will characterise all truthful mechanisms in single-parameter domains.

Single-parameter domains

Task: We will characterise all truthful mechanisms in single-parameter domains.

i.e., we will make a statement of the form: “A mechanism is truthful if any only if *it looks like this*”

Single-parameter domains

Task: We will characterise all truthful mechanisms in single-parameter domains.

i.e., we will make a statement of the form: "A mechanism is truthful if any only if *it looks like this*"

- its social choice function (allocation function) *looks like this* and

Single-parameter domains

Task: We will characterise all truthful mechanisms in single-parameter domains.

i.e., we will make a statement of the form: “A mechanism is truthful if any only if *it looks like this*”

- its social choice function (allocation function) *looks like this* and

- its payment function *looks like this.*

Monotonicity

Monotonicity

Definition (monotonicity): A social choice function f in the single-parameter domain is called monotone (in the agent's value v_i), if, for every v_{-i} and every $v_i' \geq v_i$, we have that

Monotonicity

Definition (monotonicity): A social choice function f in the single-parameter domain is called monotone (in the agent's value v_i), if, for every v_{-i} and every $v_i' \geq v_i$, we have that

$$f(v_i, v_{-i}) \in W_i \Rightarrow f(v_i', v_{-i}) \in W_i$$

Monotonicity

Definition (monotonicity): A social choice function f in the single-parameter domain is called monotone (in the agent's value v_i), if, for every v_{-i} and every $v'_i \geq v_i$, we have that

$$f(v_i, v_{-i}) \in W_i \Rightarrow f(v'_i, v_{-i}) \in W_i$$

i.e., if the value of agent i increases, then, if i was winning before, i is still winning.

Pictorially, in a single-item auction

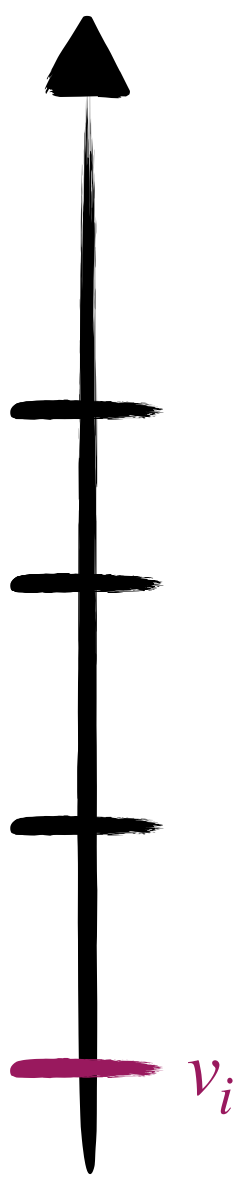

i does not win

Pictorially, in a single-item auction

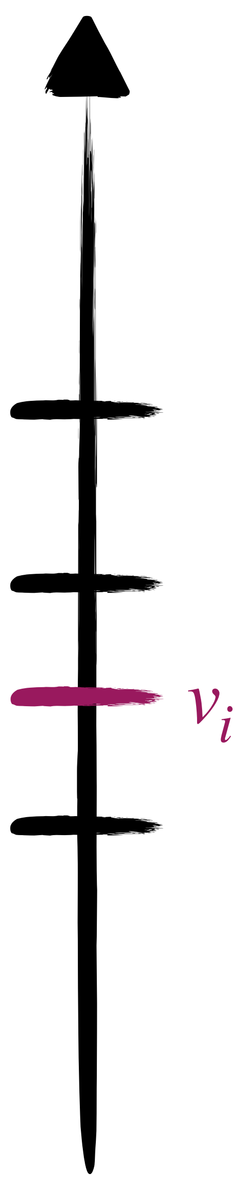

i does not win

Pictorially, in a single-item auction

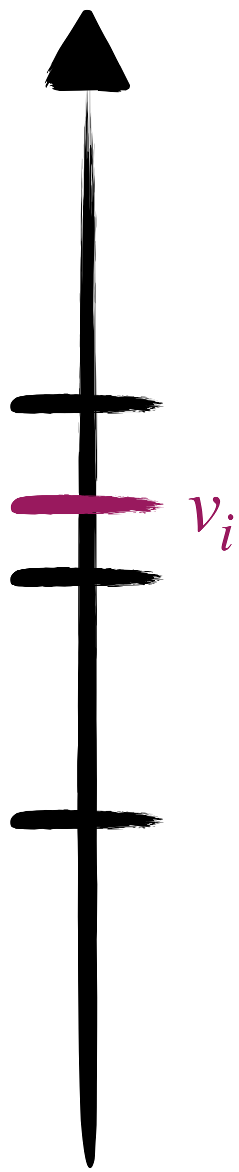

i does not win

Pictorially, in a single-item auction

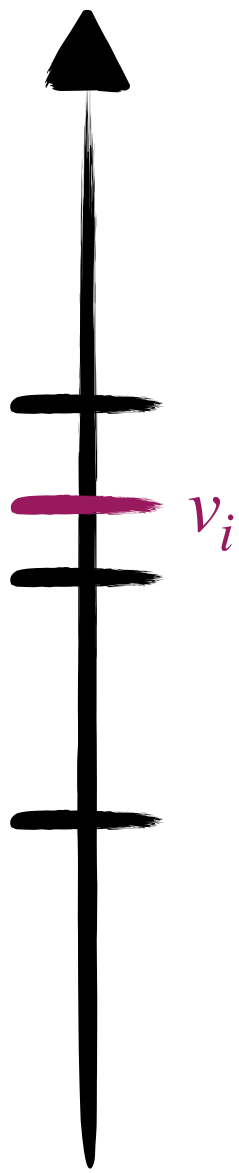

i wins

Pictorially, in a single-item auction

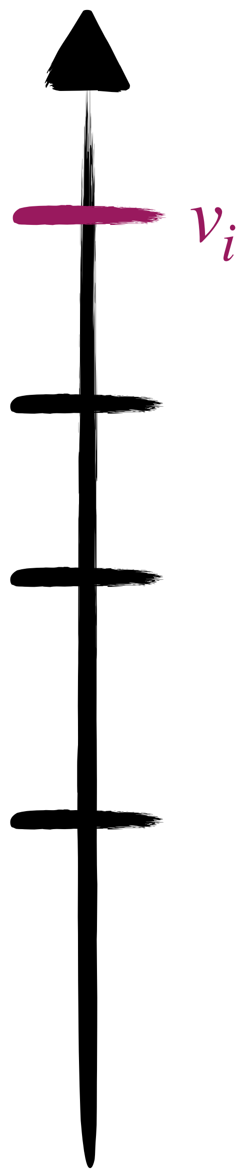

i wins

Critical value

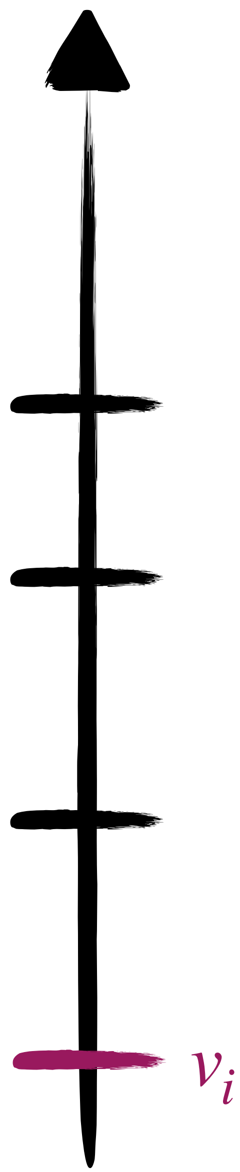

i does not win

Critical value

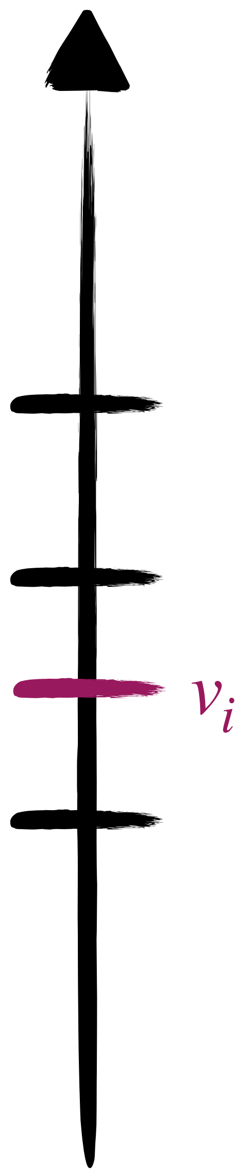

i does not win

Critical value

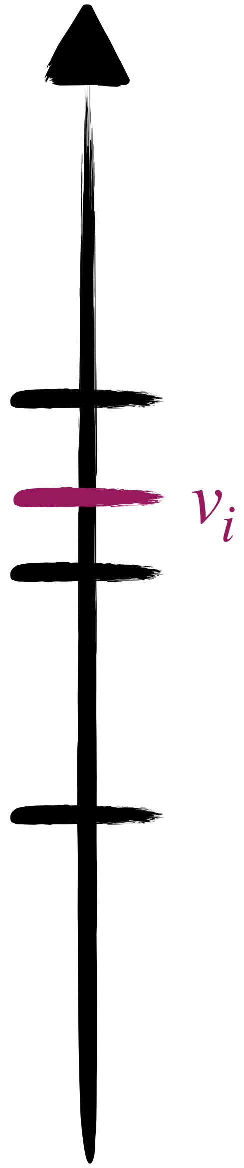

i does not win

Critical value

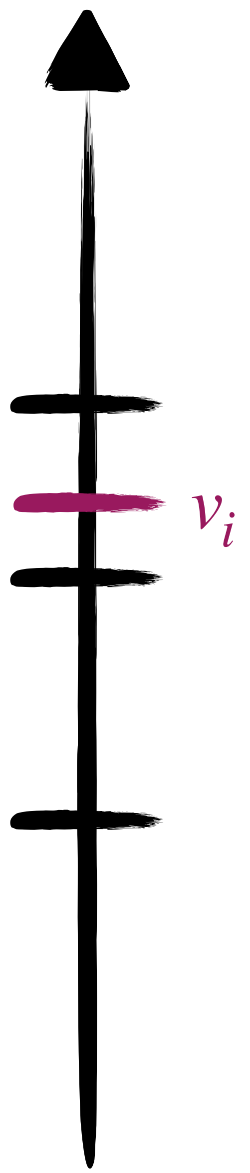

i wins

Critical value

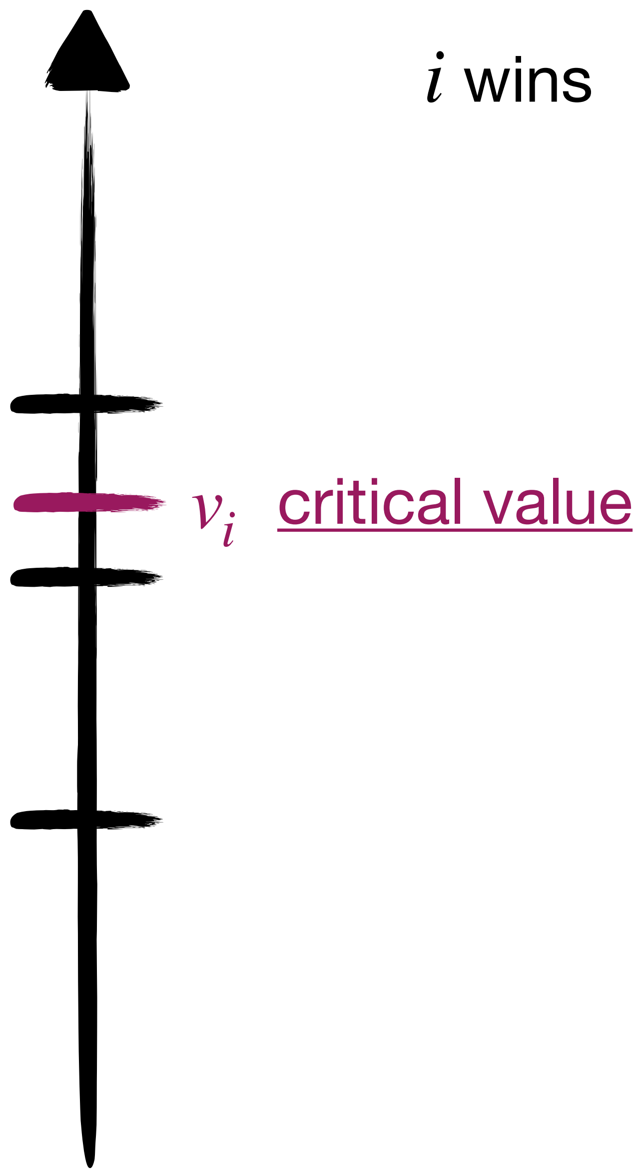

Critical value

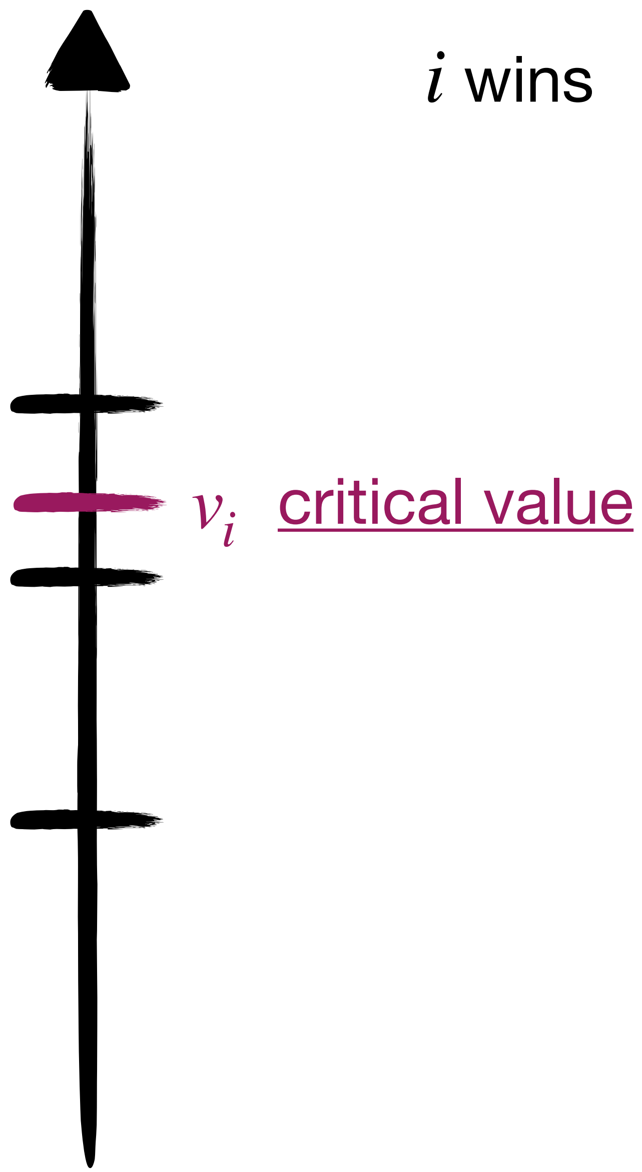

Definition (critical value): The critical value of a social choice function f in the single-parameter domain is

$$c_i(v_{-i}) = \sup_{v_i : f(v_i, v_{-i}) \neq W_i} v_i$$

Second-price auctions (SPA)

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

> If there are multiple such bidders, one is chosen at random.

The winner needs to pay the bid of the second highest bidder, all other bidders do not pay anything.

William Vickrey

Second-price auctions (SPA)

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

> If there are multiple such bidders, one is chosen at random.

The winner needs to pay the bid of the second highest bidder, all other bidders do not pay anything.

William Vickrey

What is the critical value in the SPA?

Myerson's Characterisation

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981): Let $(f, p_1, \ldots, p_n)$ be a mechanism on a single-parameter domain, *for which losers pay* 0. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981): Let $(f, p_1, \ldots, p_n)$ be a mechanism on a single-parameter domain, *for which losers pay* 0. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): f is *monotone*.

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981): Let $(f, p_1, \ldots, p_n)$ be a mechanism on a single-parameter domain, *for which losers pay* 0. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): f is *monotone*.

(2) Condition on the payments: The payment p_i of every winner is the *critical value*.

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981): Let $(f, p_1, \ldots, p_n)$ be a mechanism on a single-parameter domain, *for which losers pay* 0. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): f is *monotone*.

(2) Condition on the payments: The payment p_i of every winner is the *critical value*.

Formally, for every i, v_i, and v_{-i} such that $f(v_i, v_{-i}) \in W_i$, we have that $p_i = c_i(v_{-i})$.

Second-price auctions (SPA)

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

> If there are multiple such bidders, one is chosen at random.

The winner needs to pay the bid of the second highest bidder, all other bidders do not pay anything.

William Vickrey

Second-price auctions (SPA)

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

If there are multiple such bidders, one is chosen at random.

The winner needs to pay the bid of the second highest bidder, all other bidders do not pay anything.

William Vickrey

What is the critical value in the SPA?

Second-price auctions (SPA)

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

If there are multiple such bidders, one is chosen at random.

The winner needs to pay the bid of the second highest bidder, all other bidders do not pay anything.

William Vickrey

What is the critical value in the SPA? What is the payment?

Second-price auctions with reserve

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

If there are multiple such bidders, one is chosen at random.

The winner needs to pay the maximum of the bid of the second highest bidder and a reserve price $\hat{p}$, all other bidders do not pay anything.

Second-price auctions with reserve

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

If there are multiple such bidders, one is chosen at random.

The winner needs to pay the maximum of the bid of the second highest bidder and a reserve price $\hat{p}$, all other bidders do not pay anything.

What is the critical value in the SPA with reserve?

Second-price auctions with reserve

These are also *sealed-bid auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

If there are multiple such bidders, one is chosen at random.

The winner needs to pay the maximum of the bid of the second highest bidder and a reserve price $\hat{p}$, all other bidders do not pay anything.

What is the critical value in the SPA with reserve? What is the payment?

Why is the FPA not truthful then?

Possible reason: The SCF (allocation) is not monotone. Is it?

Possible reason: The payment is not the critical value.

Critical value

i does not win

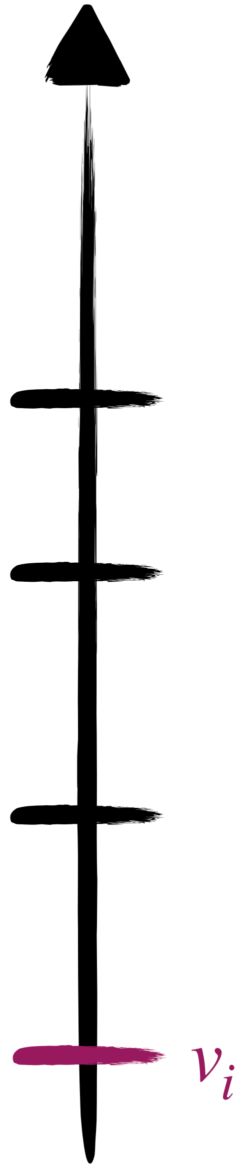

Critical value

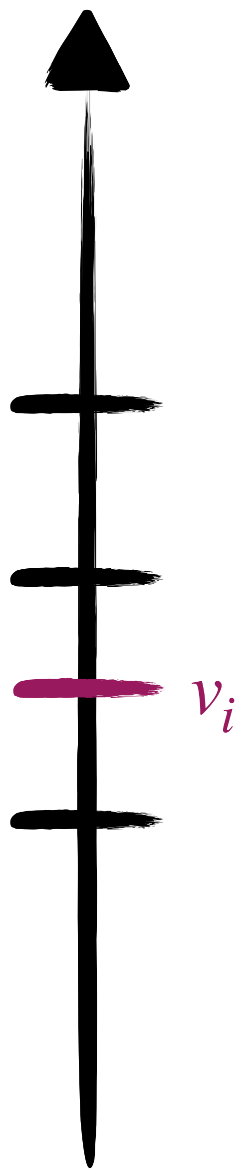

i does not win

Critical value

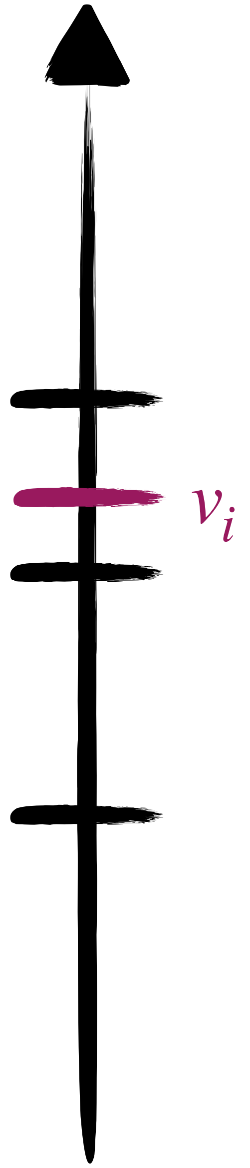

i does not win

Critical value

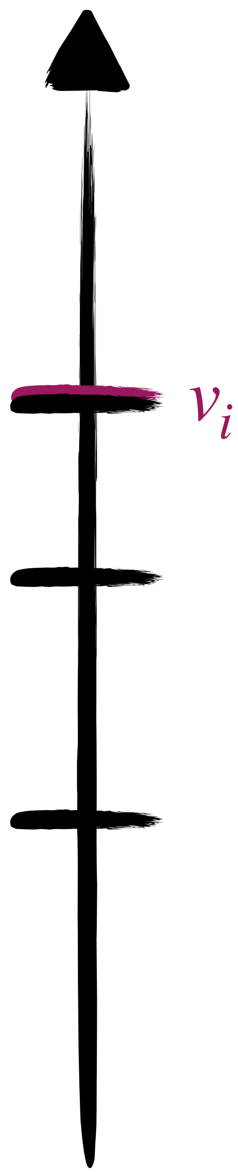

i does not win

Critical value

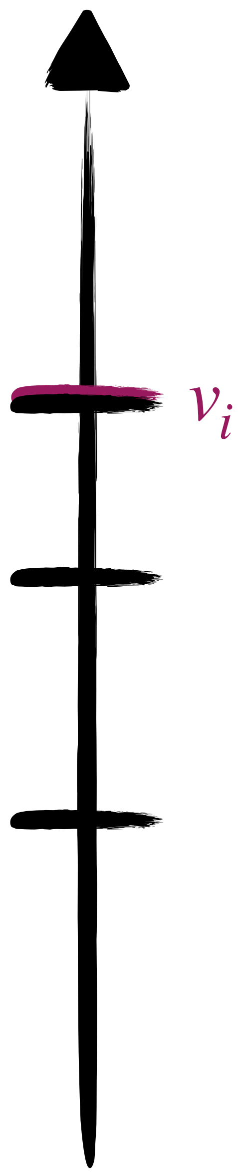

i wins

Critical value

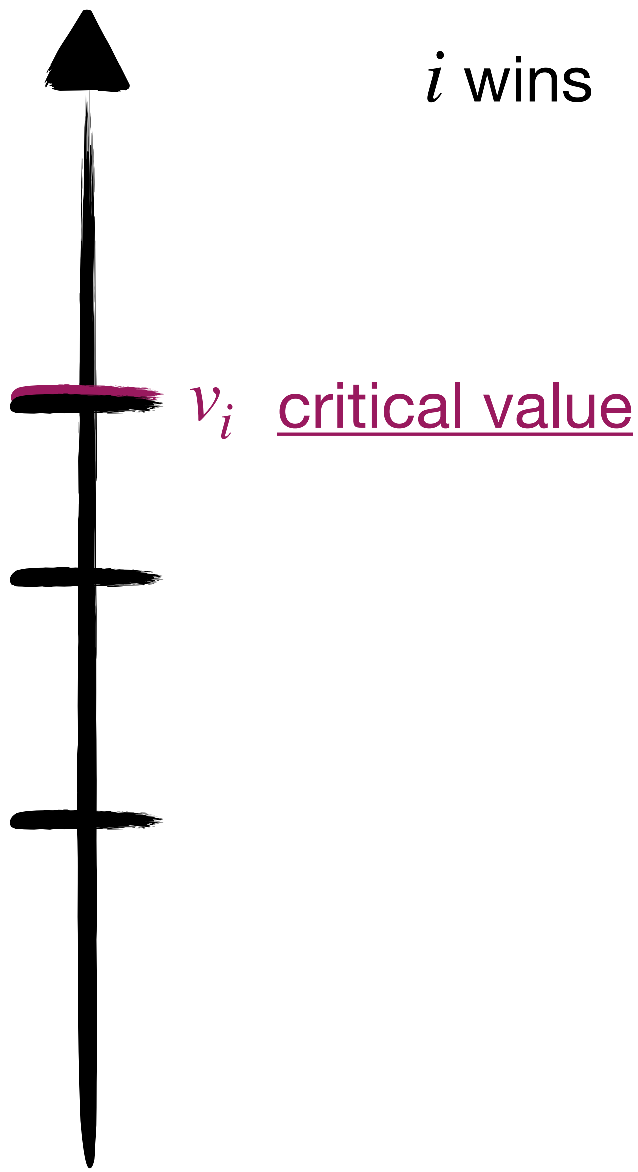

Critical value

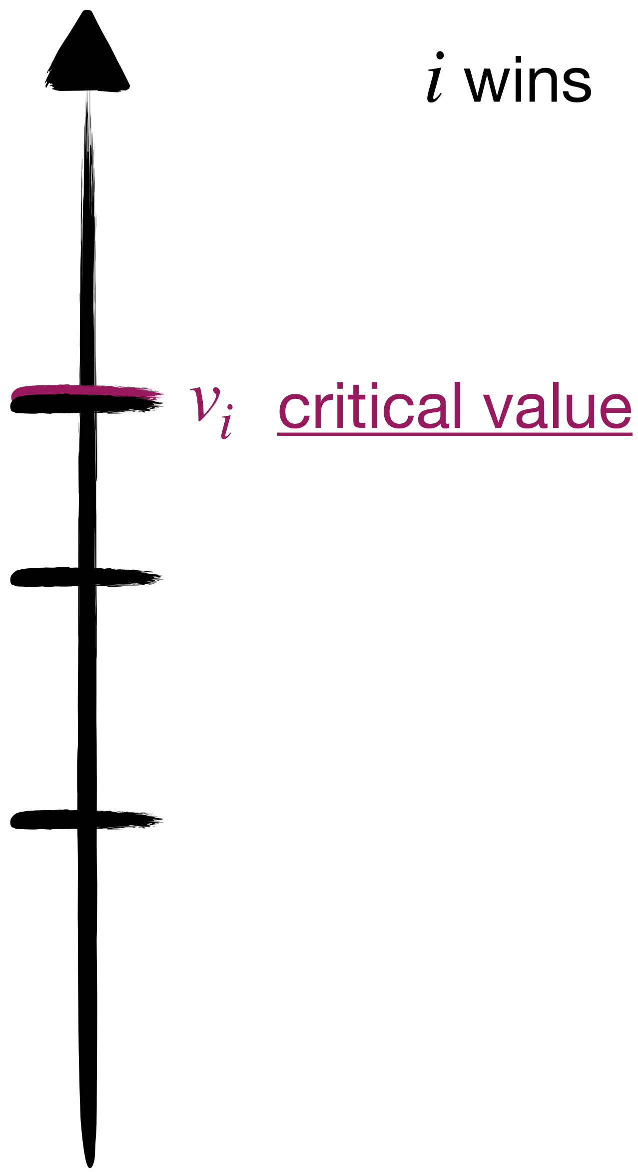

The payment is not the critical value.

Critical value

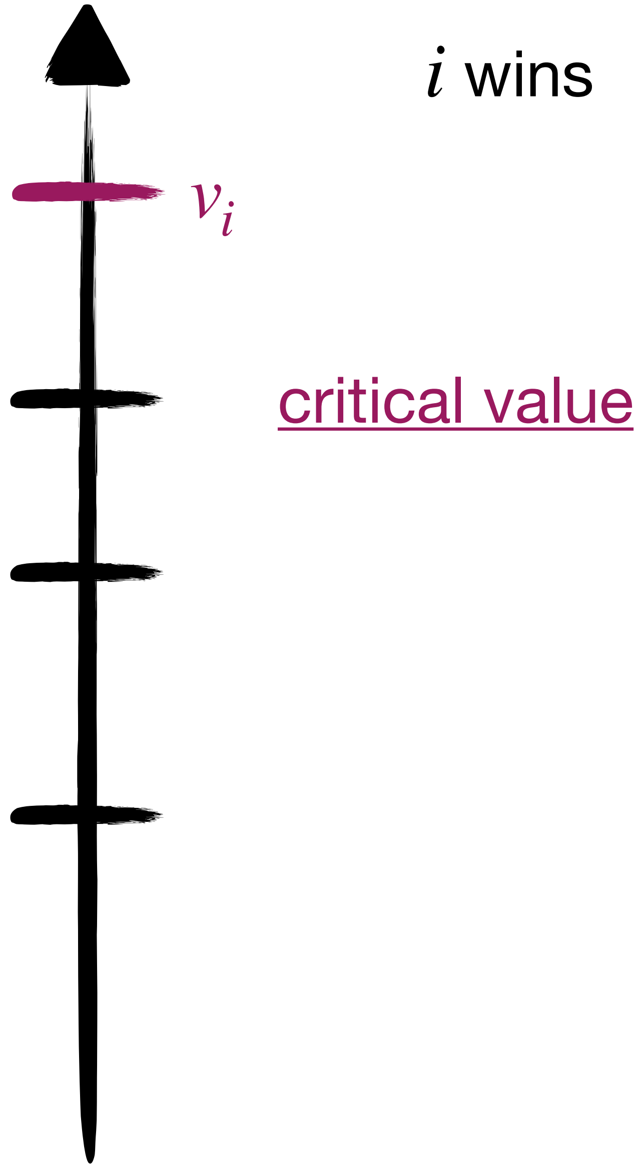

The payment is not the critical value.

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981): Let $(f, p_1, \ldots, p_n)$ be a mechanism on a single-parameter domain, *for which losers pay* 0. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): f is *monotone*.

(2) Condition on the payments: The payment p_i of every winner is the *critical value*.

Formally, for every i, v_i, and v_{-i} such that $f(v_i, v_{-i}) \in W_i$, we have that $p_i = c_i(v_{-i})$.

Proof: Monotone + Critical Value Payment ⇒ Truthful

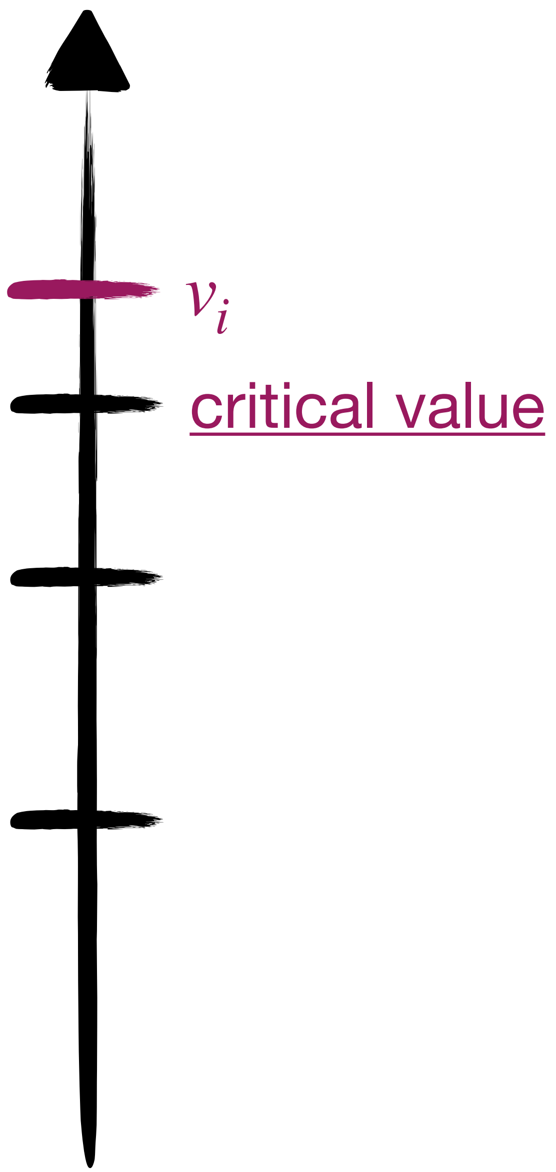

Proof: Monotone + Critical Value Payment ⇒ Truthful

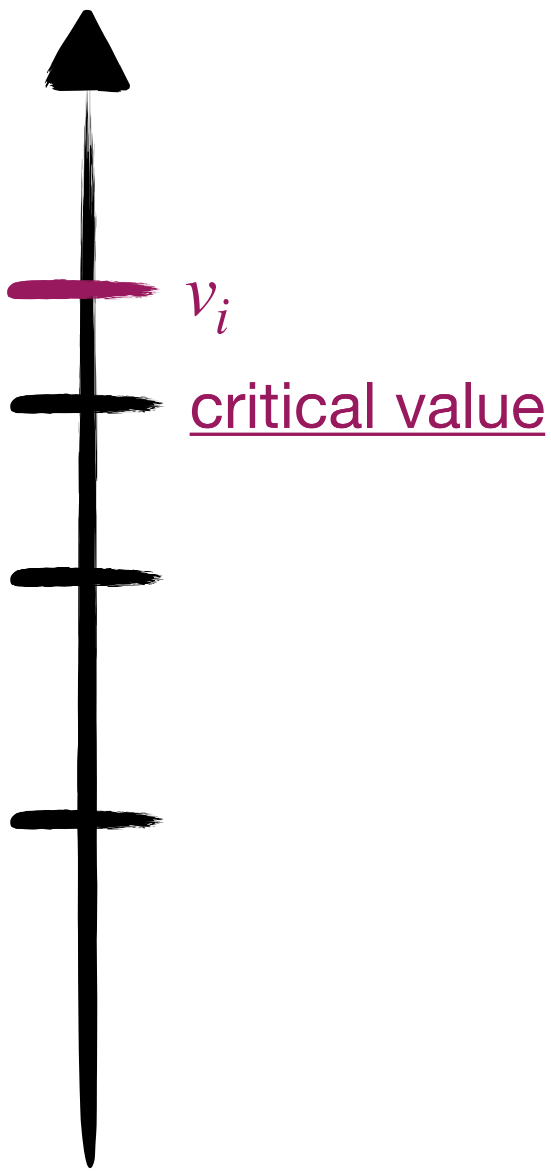

Assume with real value v_i that agent i is winning.

Proof: Monotone + Critical Value Payment ⇒ Truthful

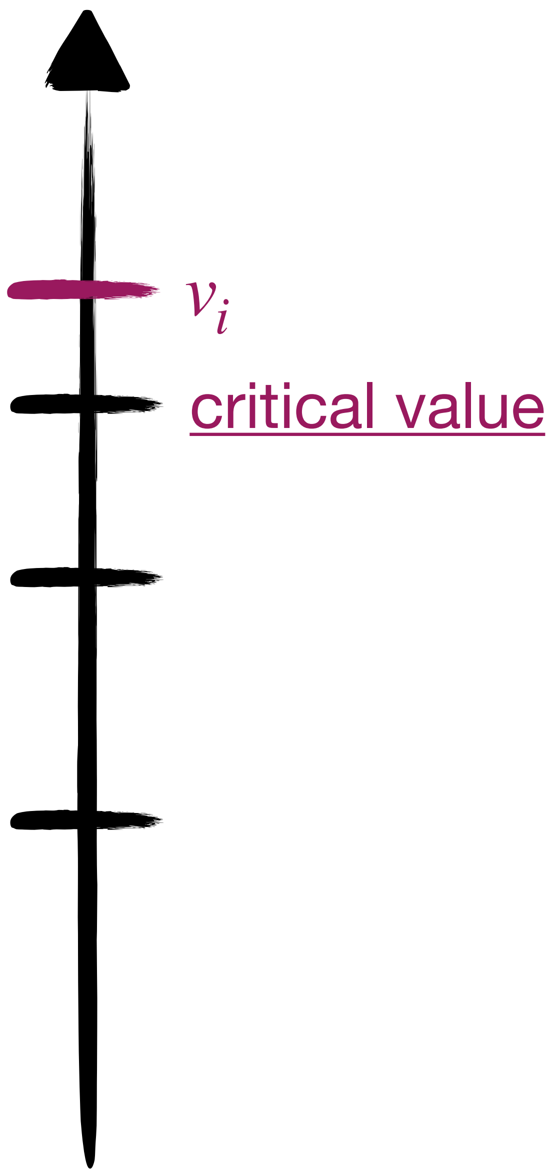

Assume with real value v_i that agent i is winning.

Utility is is $v_i - p_i = v_i - c_i(v_{-i})$.

Proof: Monotone + Critical Value Payment ⇒ Truthful

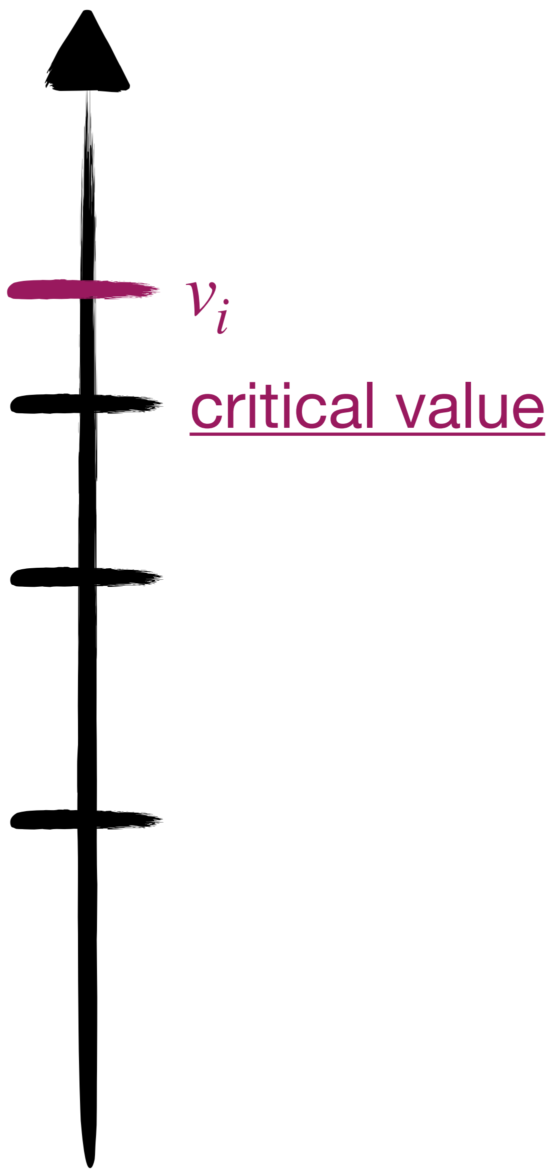

Assume with real value v_i that agent i is winning.

Utility is is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at least 0, why?

Proof: Monotone + Critical Value Payment ⇒ Truthful

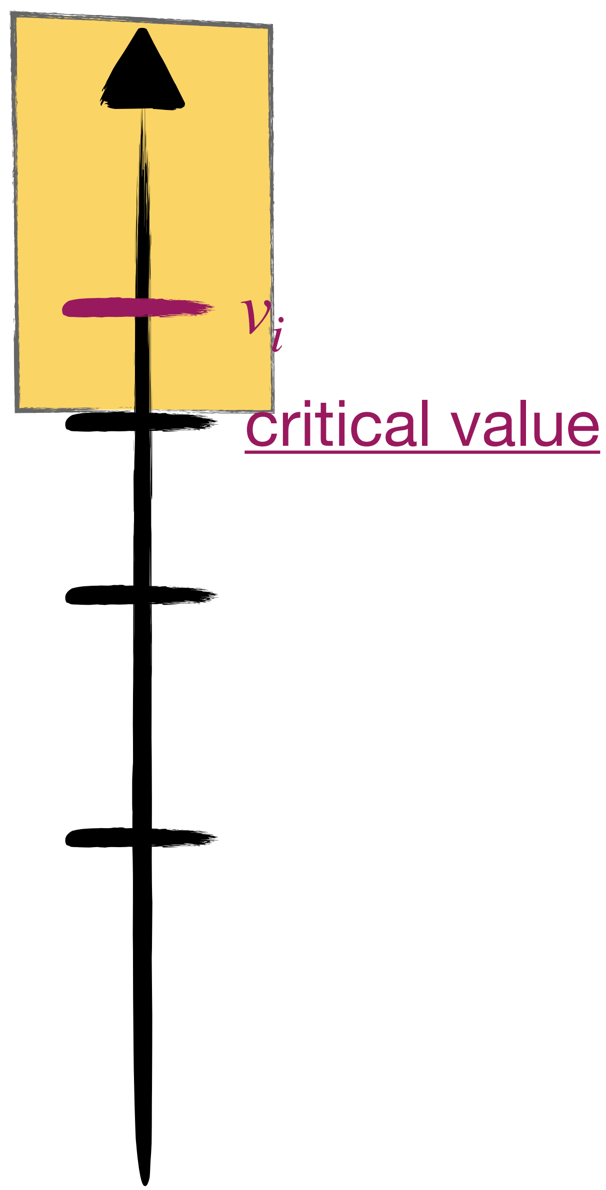

Assume with real value v_i that agent i is winning.

Utility is is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at least 0, why?

Proof: Monotone + Critical Value Payment ⇒ Truthful

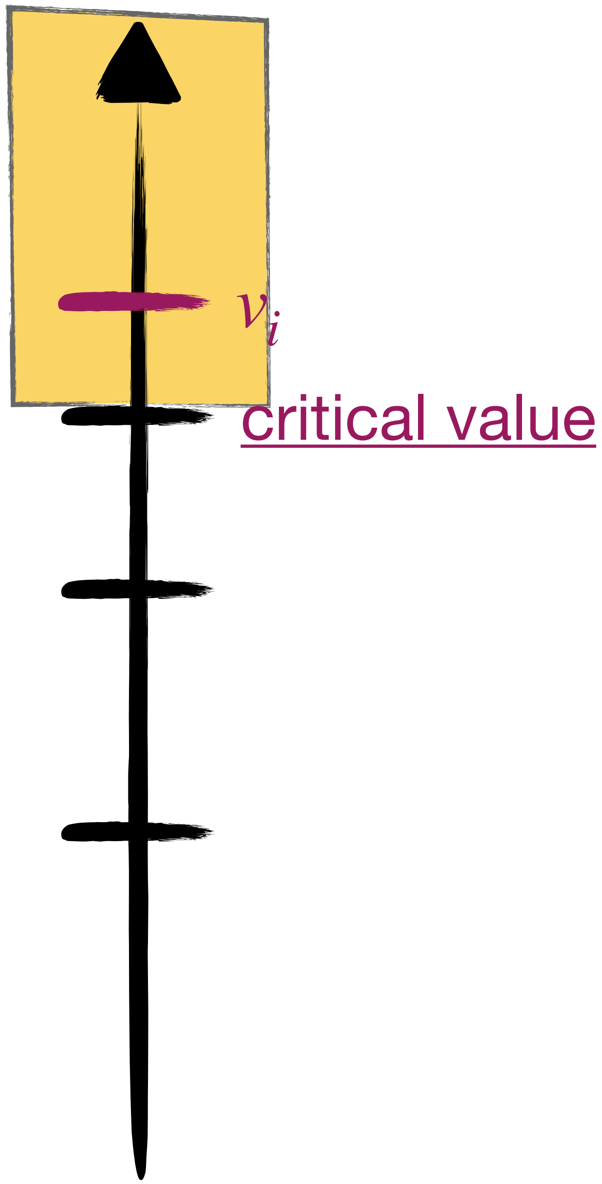

Assume with real value v_i that agent i is winning.

Utility is is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at least 0, why?

Still winning, still paying the same, same utility.

Proof: Monotone + Critical Value Payment ⇒ Truthful

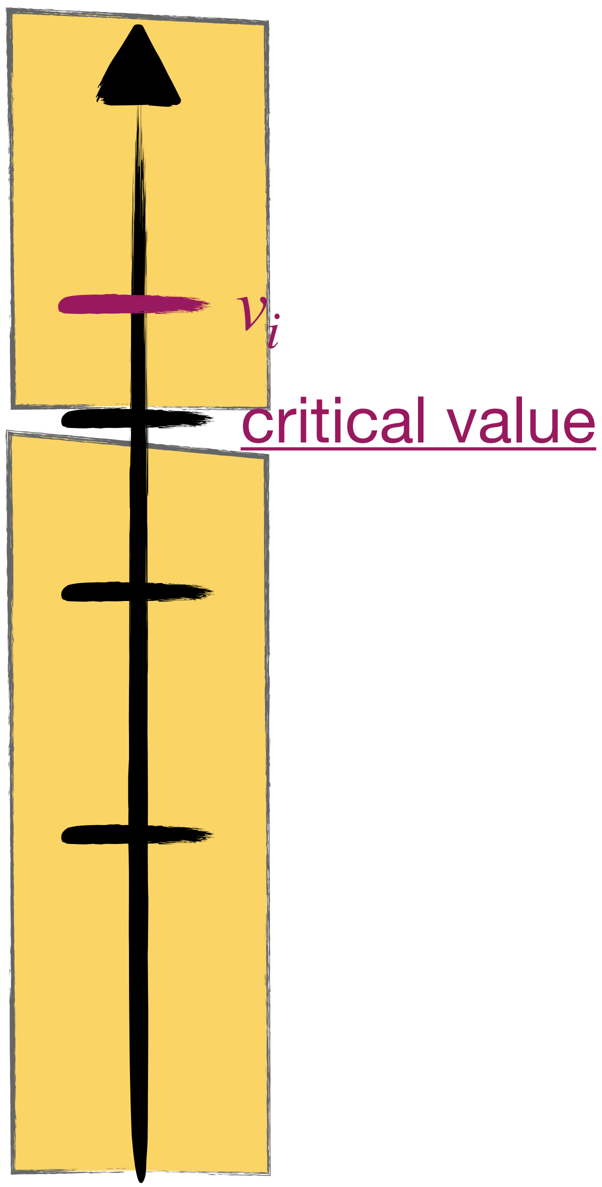

Assume with real value v_i that agent i is winning.

Utility is is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at least 0, why?

Still winning, still paying the same, same utility.

Proof: Monotone + Critical Value Payment ⇒ Truthful

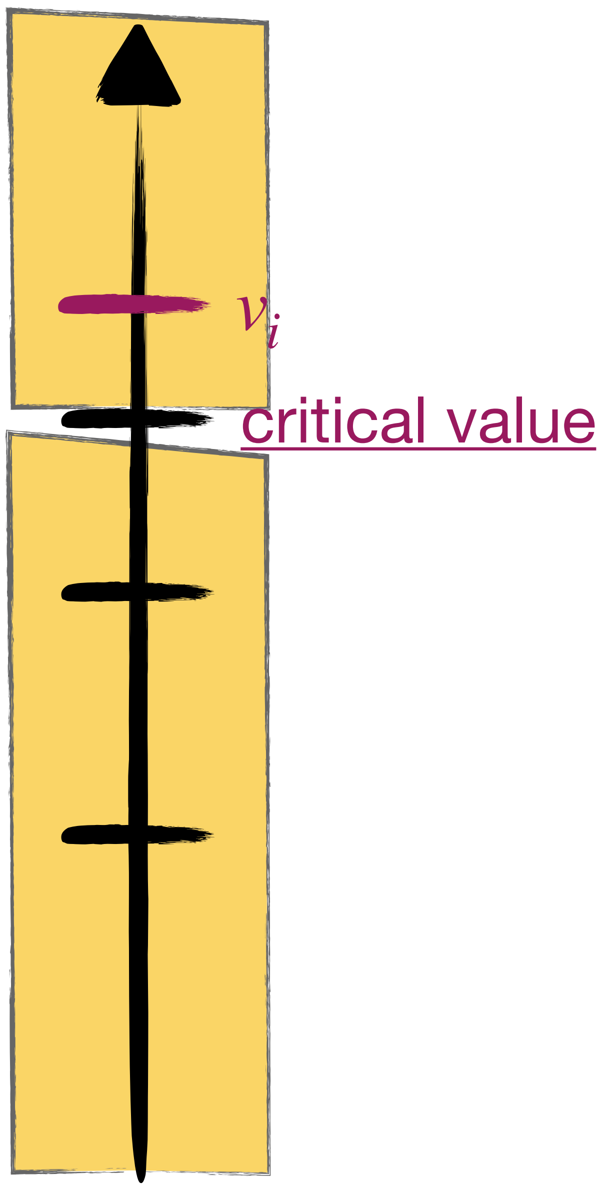

Assume with real value v_i that agent i is winning.

Utility is is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at least 0, why?

Still winning, still paying the same, same utility.

May be losing, then the utility is 0.

Proof: Monotone + Critical Value Payment ⇒ Truthful

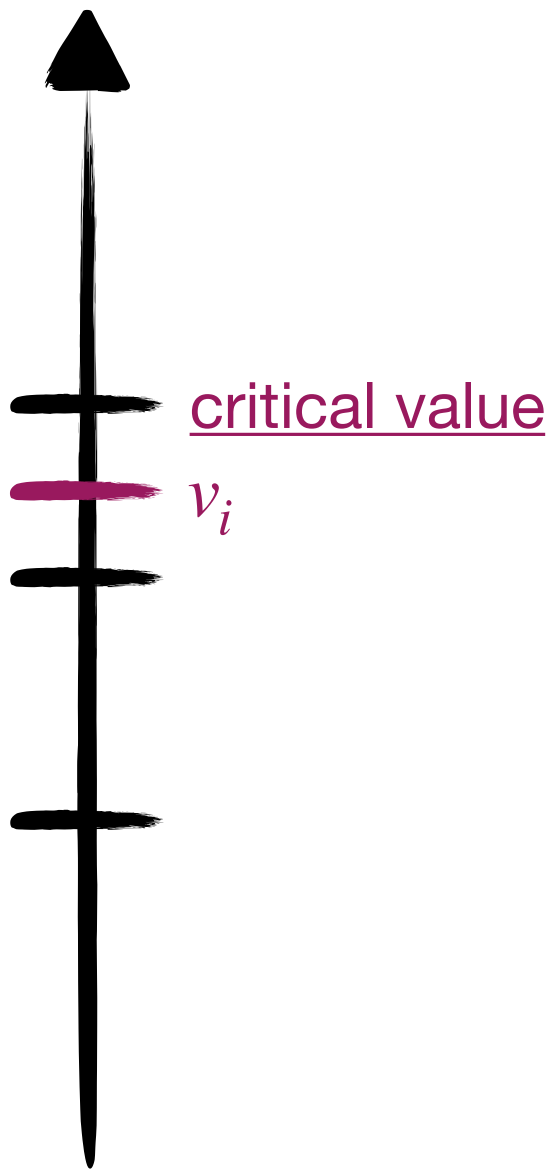

Proof: Monotone + Critical Value Payment ⇒ Truthful

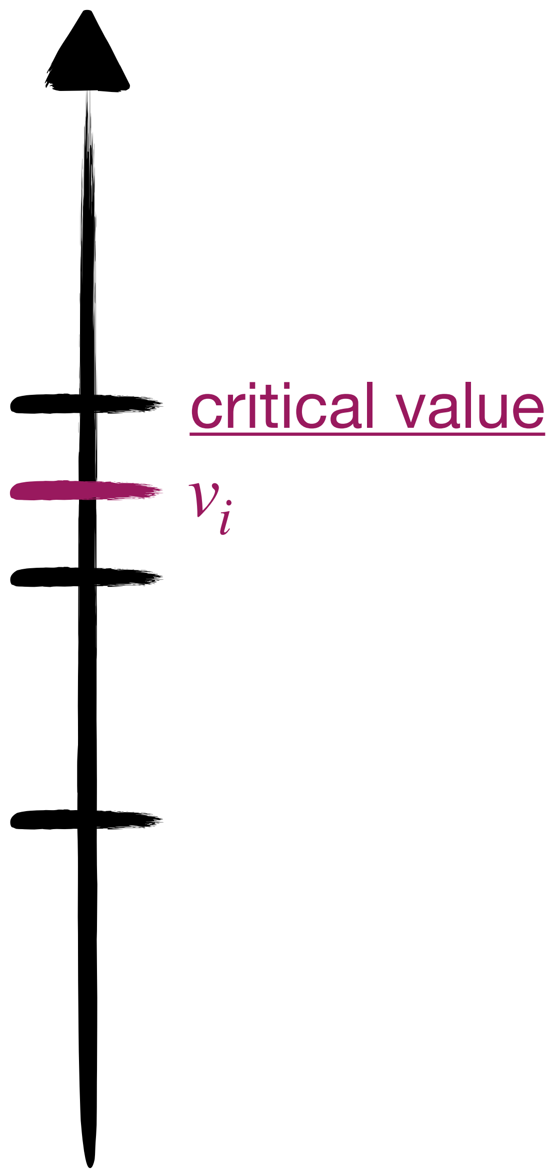

Assume with real value v_i that agent i is losing.

Proof: Monotone + Critical Value Payment ⇒ Truthful

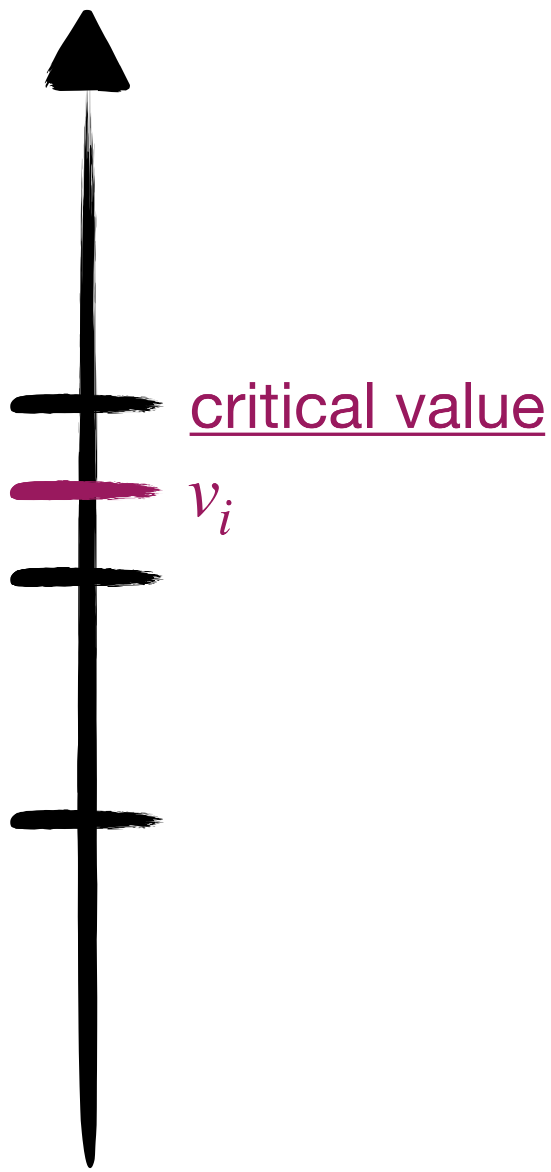

Assume with real value v_i that agent i is losing.

Utility of winning is $v_i - p_i = v_i - c_i(v_{-i})$.

Proof: Monotone + Critical Value Payment ⇒ Truthful

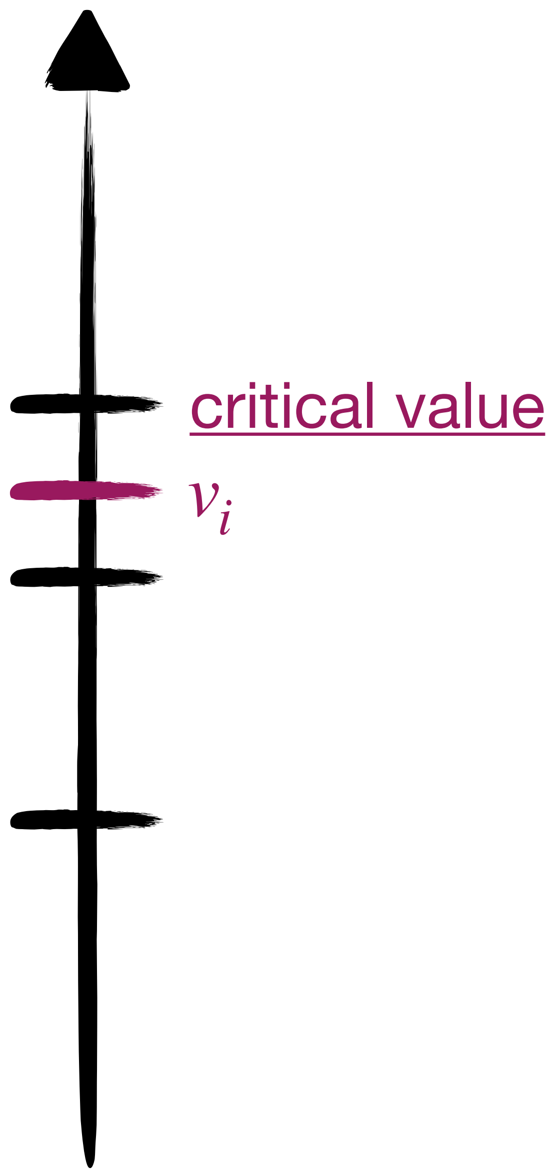

Assume with real value v_i that agent i is losing.

Utility of winning is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at most 0, why?

Proof: Monotone + Critical Value Payment ⇒ Truthful

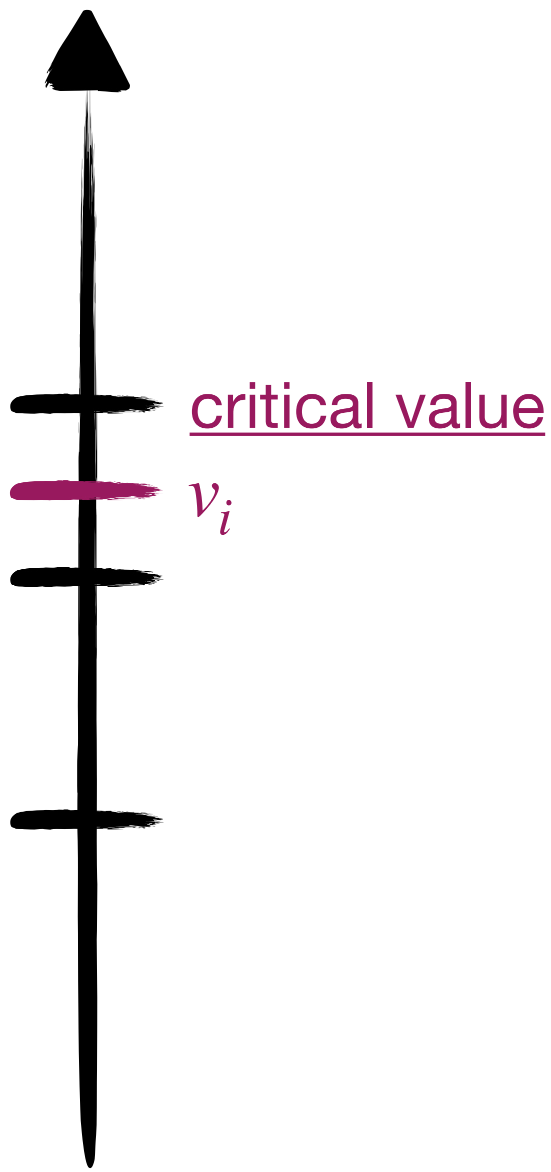

Assume with real value v_i that agent i is losing.

Utility of winning is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at most 0, why?

Losing does not change the payment, winning is not better.

Proof: Monotone + Critical Value Payment ⇒ Truthful

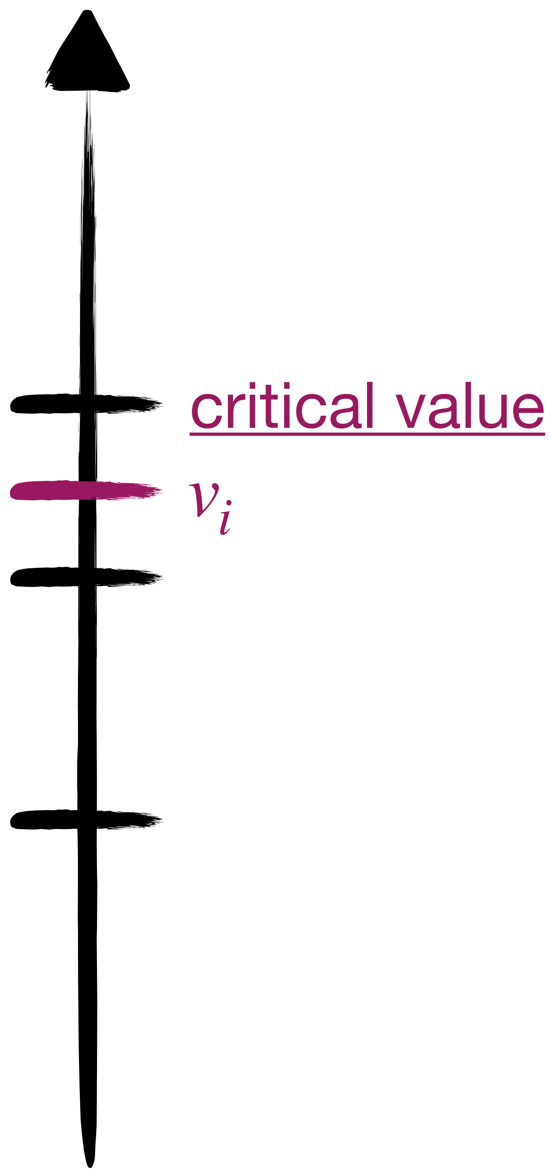

Assume with real value v_i that agent i is losing.

Utility of winning is $v_i - p_i = v_i - c_i(v_{-i})$.

This utility is at most 0, why?

Losing does not change the payment, winning is not better.

Where was monotonicity used really?

Proof: Truthful $\Rightarrow$ Monotone + Critical Value Payment

Proof: Truthful $\Rightarrow$ Monotone + Critical Value Payment

Monotonicity: Assume by contradiction that f is not monotone.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Monotonicity: Assume by contradiction that f is not monotone.

That means that with a higher value $v_i' > v_i$ the agent *loses*, whereas with v_i the agent *wins*. The payment in the latter case is $p_i(v_i, v_{-i})$.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Monotonicity: Assume by contradiction that f is not monotone.

That means that with a higher value $v_i' > v_i$ the agent *loses*, whereas with v_i the agent *wins*. The payment in the latter case is $p_i(v_i, v_{-i})$.

By truthfulness, $u_i(v_i, p_i) = v_i - p(v_i, v_{-i}) \geq 0$, as otherwise the agent with real value v_i would have an incentive to *misreport* v_i', lose, and get a utility of 0.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Monotonicity: Assume by contradiction that f is not monotone.

That means that with a higher value $v_i' > v_i$ the agent *loses*, whereas with v_i the agent *wins*. The payment in the latter case is $p_i(v_i, v_{-i})$.

By truthfulness, $u_i(v_i, p_i) = v_i - p(v_i, v_{-i}) \geq 0$, as otherwise the agent with real value v_i would have an incentive to *misreport* v_i', lose, and get a utility of 0.

By truthfulness, $u_i(v_i', p_i) = v_i' - p(v_i, v_{-i}) \leq 0$, as otherwise the agent with real value v_i' would have an incentive to *misreport* v_i, win, and get a positive utility.

Proof: Truthful $\Rightarrow$ Monotone + Critical Value Payment

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Proof: Truthful $\Rightarrow$ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

Proof: Truthful $\Rightarrow$ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

(1) $p > c_i(v_{-i})$. Let $v_i' > c_i(v_{-i})$ and $v_i' < p$.

Proof: Truthful $\Rightarrow$ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

(1) $p > c_i(v_{-i})$. Let $v_i' > c_i(v_{-i})$ and $v_i' < p$.

v_i' is a winning bid.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

(1) $p > c_i(v_{-i})$. Let $v_i' > c_i(v_{-i})$ and $v_i' < p$.

v_i' is a winning bid.

By the Claim above, $p(v_i', v_{-i}) = p(v_i, v_{-i})$.

Proof: Truthful $\Rightarrow$ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

(1) $p > c_i(v_{-i})$. Let $v_i' > c_i(v_{-i})$ and $v_i' < p$.

v_i' is a winning bid.

By the Claim above, $p(v_i', v_{-i}) = p(v_i, v_{-i})$.

An agent with true value v_i' now has negative utility, so it would prefer to bid 0 and lose, violating truthfulness.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

(2) $p < c_i(v_{-i})$. Let $v_i' < c_i(v_{-i})$ and $v_i' > p$.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

(2) $p < c_i(v_{-i})$. Let $v_i' < c_i(v_{-i})$ and $v_i' > p$.

v_i' is a losing bid.

Proof: Truthful ⇒ Monotone + Critical Value Payment

Payment: First we prove the following claim:

Claim: Fix any v_{-i}. Let v_i and v_i' be such that bidder i wins with both. Then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.

Proof of claim: Otherwise an agent with true value v_i or an agent with true value v_i' could increase its utility by misreporting the other value.

Now assume by contradiction that some winning agent pays

(2) $p < c_i(v_{-i})$. Let $v_i' < c_i(v_{-i})$ and $v_i' > p$.

v_i' is a losing bid.

An agent with true value v_i' has 0 utility, and it would prefer to bid v_i and win, gaining positive utility, and violating truthfulness.

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981): Let $(f, p_1, \ldots, p_n)$ be a mechanism on a single-parameter domain, *for which losers pay* 0. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): f is *monotone*.

(2) Condition on the payments: The payment p_i of every winner is the *critical value*.

Formally, for every i, v_i, and v_{-i} such that $f(v_i, v_{-i}) \in W_i$, we have that $p_i = c_i(v_{-i})$.

Randomized Mechanisms

Randomized Mechanisms

So far we have been talking about *deterministic* mechanisms, in which an agent is either a winning agent or a losing agent.

Randomized Mechanisms

So far we have been talking about *deterministic* mechanisms, in which an agent is either a winning agent or a losing agent.

We can also have *randomised* mechanisms, in which the agent is a winning agent with some probability.

Randomized Mechanisms

So far we have been talking about *deterministic* mechanisms, in which an agent is either a winning agent or a losing agent.

We can also have *randomised* mechanisms, in which the agent is a winning agent with some probability.

e.g., in single-item auctions, the agent wins the item with some probability $w_i(v_i, v_{-i})$.

Randomized Mechanisms

So far we have been talking about *deterministic* mechanisms, in which an agent is either a winning agent or a losing agent.

We can also have *randomised* mechanisms, in which the agent is a winning agent with some probability.

e.g., in single-item auctions, the agent wins the item with some probability $w_i(v_i, v_{-i})$.

Myerson's characterisation can be generalised for these mechanisms as well!

Randomized Mechanisms

Randomized Mechanisms

The utility of an agent is given by $v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$.

Randomized Mechanisms

The utility of an agent is given by $v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$.

We will consider *normalised* mechanisms in which the lowest v_i has 0 probability of winning, i.e., $w_i(v_i^\ell, v_{-i}) = 0$ for $v_i^\ell = \min_i v_i$ and incurs 0 payment, i.e., $p_i(v_i^\ell, v_{-i}) = 0$.

Myerson's Characterisation

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981):
Let $(f, p_1, \ldots, p_n)$ be a *normalised* randomised mechanism on a single-parameter domain. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981):
Let $(f, p_1, \ldots, p_n)$ be a *normalised* randomised mechanism on a single-parameter domain. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): the function $w_i(v_i, v_{-i})$ is monotonically non-decreasing in v_i.

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981):
Let $(f, p_1, \ldots, p_n)$ be a *normalised* randomised mechanism on a single-parameter domain. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): the function $w_i(v_i, v_{-i})$ is monotonically non-decreasing in v_i.

(2) Condition on the payments: The payment p_i of every winner is given by the following formula:

Myerson's Characterisation

Theorem (Myerson's Characterisation or Myerson's Lemma, Myerson 1981):
Let $(f, p_1, \ldots, p_n)$ be a *normalised* randomised mechanism on a single-parameter domain. Then, $(f, p_1, \ldots, p_n)$ is truthful if any only if the following conditions hold:

(1) Condition on the SCF (allocation): the function $w_i(v_i, v_{-i})$ is monotonically non-decreasing in v_i.

(2) Condition on the payments: The payment p_i of every winner is given by the following formula:

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^{\ell}}^{v_i} w(t, v_{-i})dt$$

Some pictures

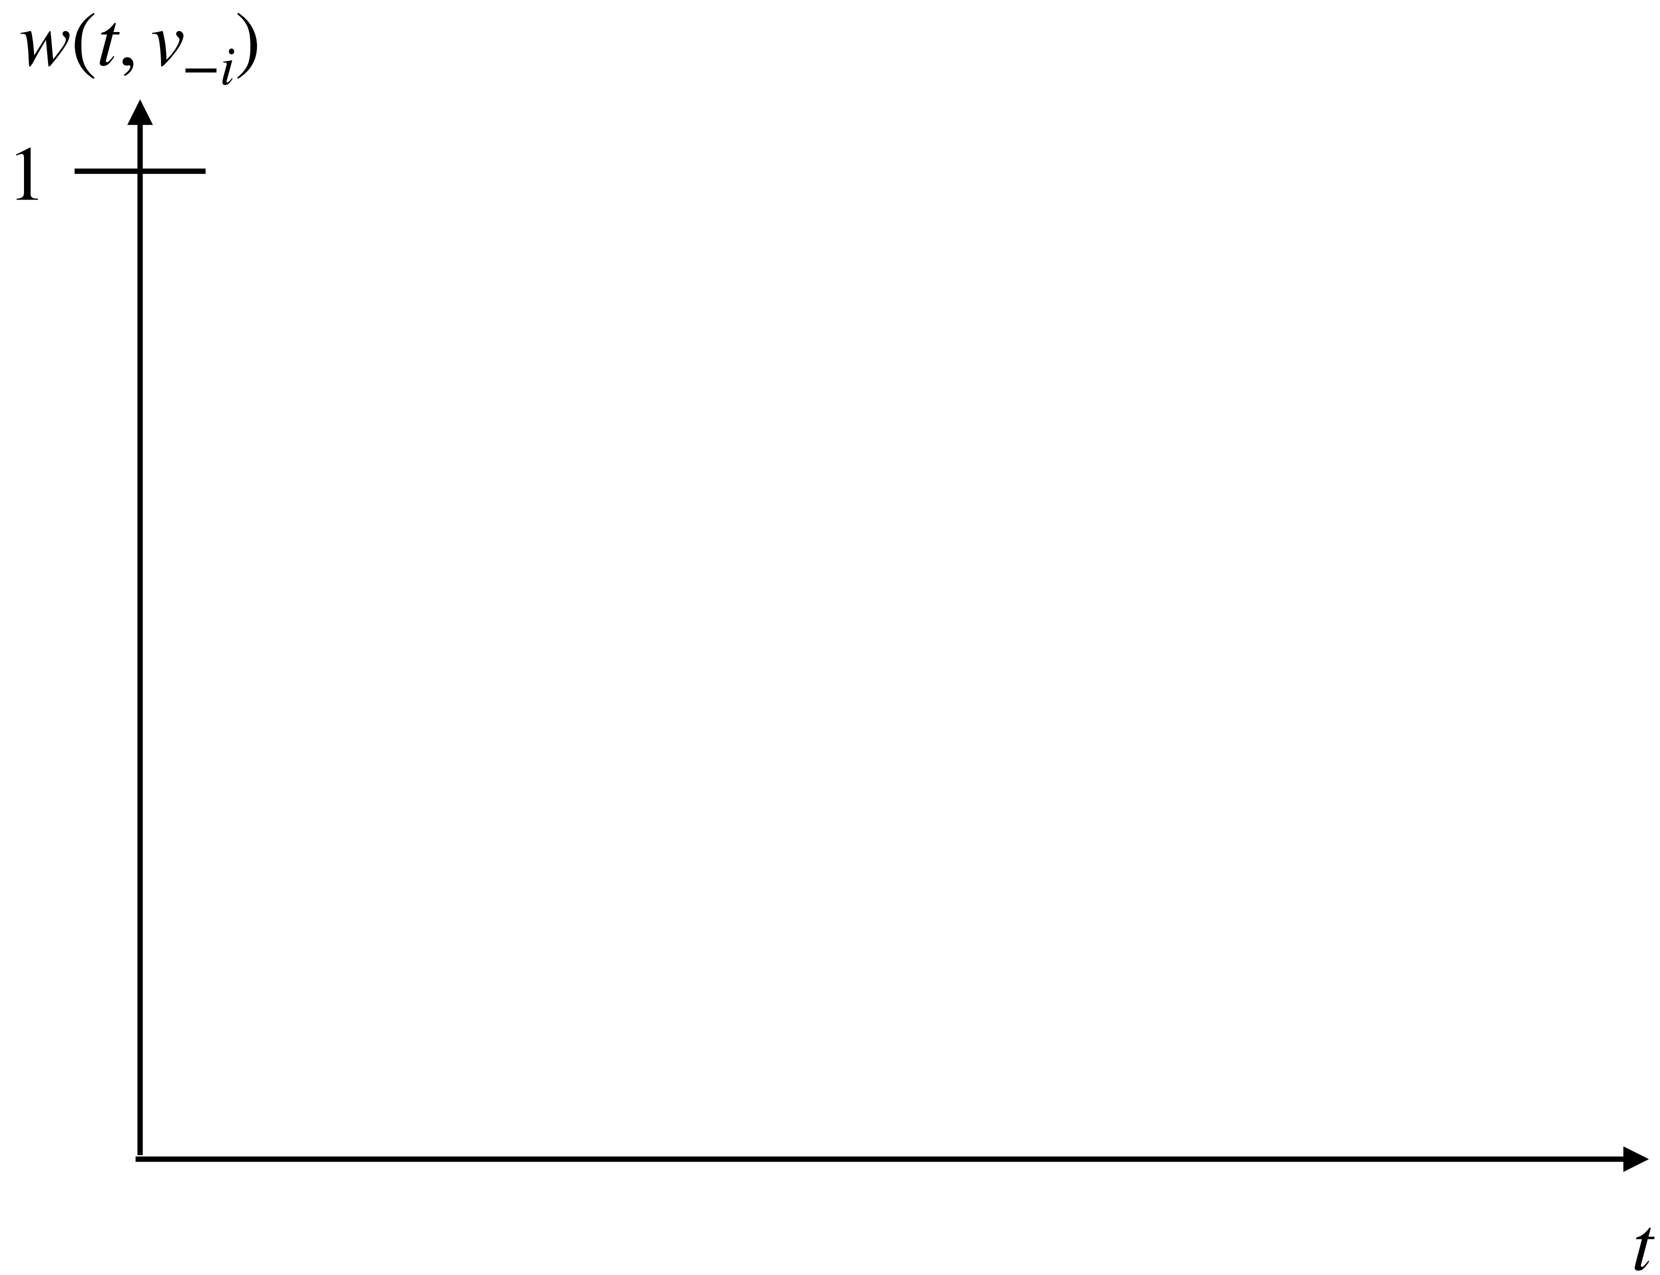

Some pictures

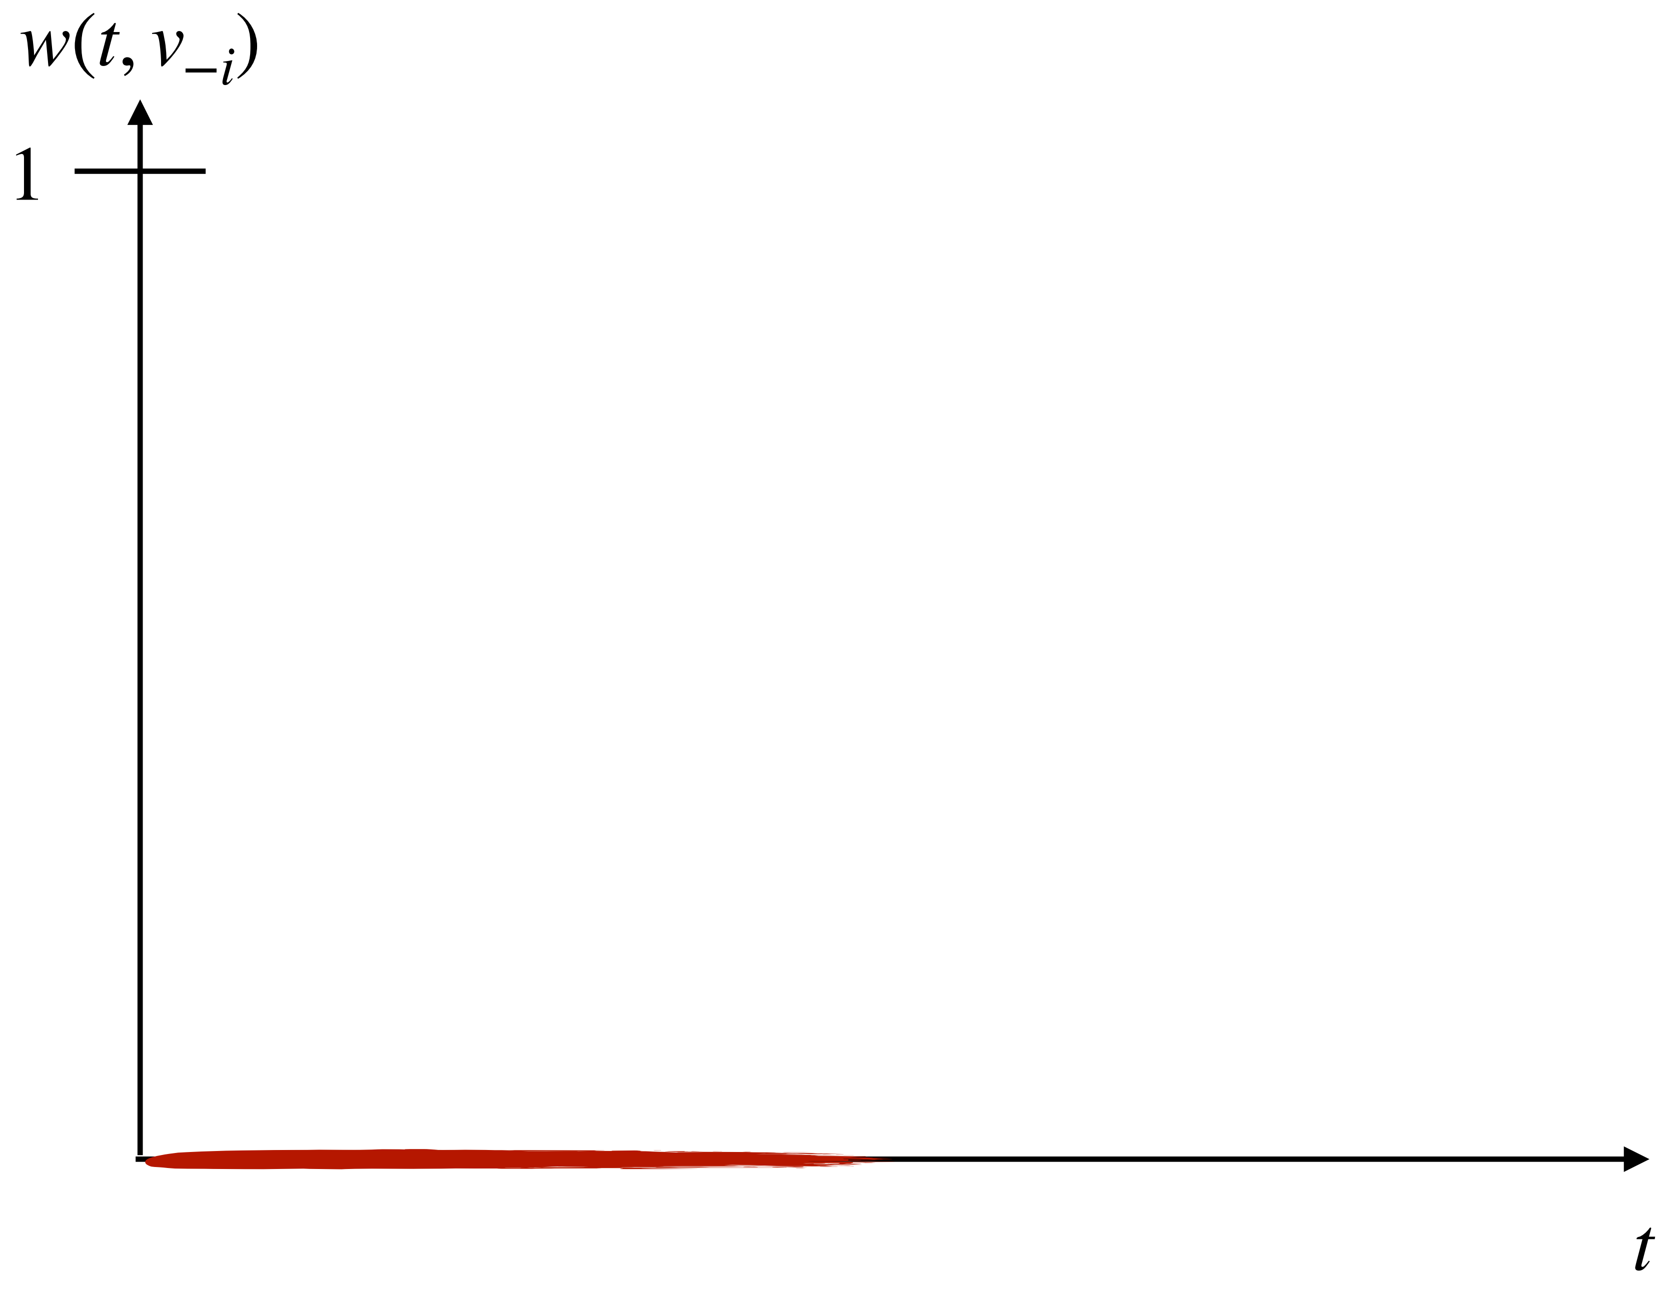

Some pictures

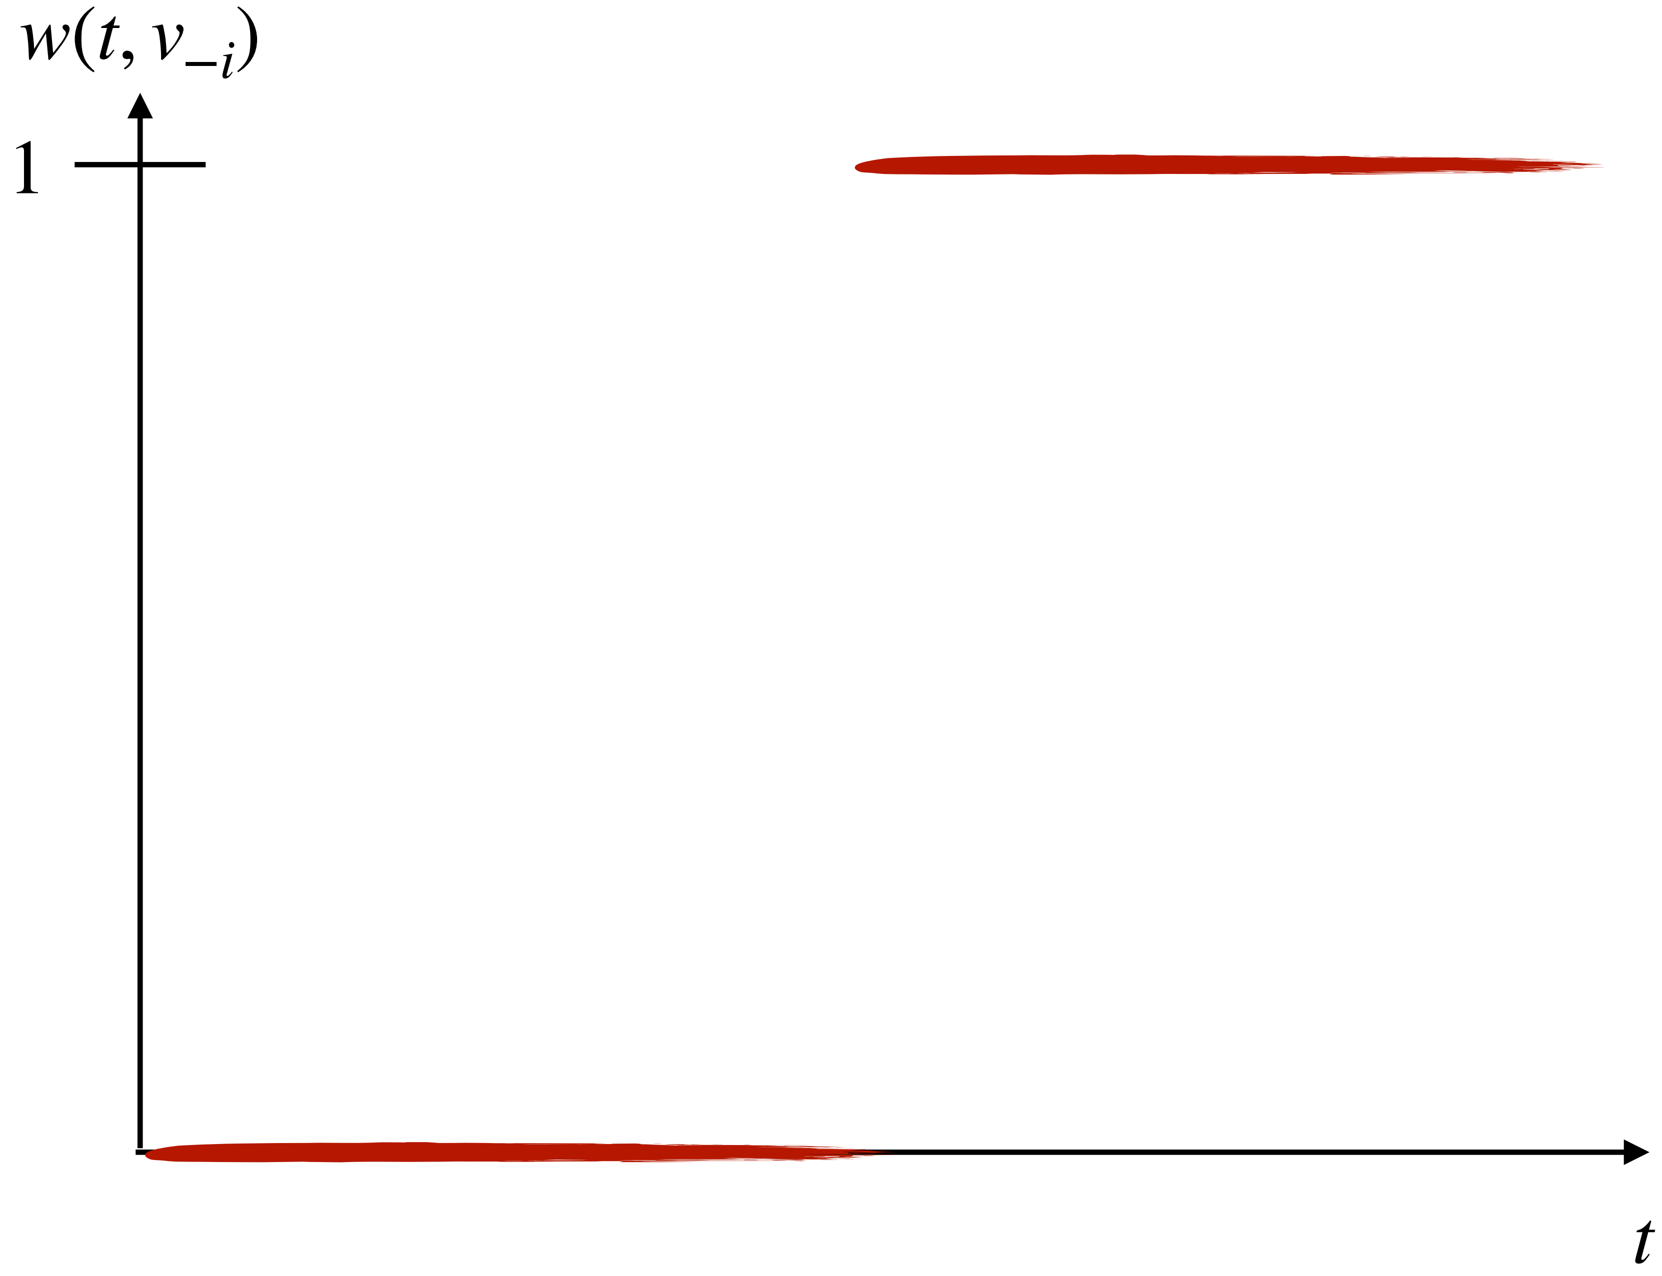

Some pictures

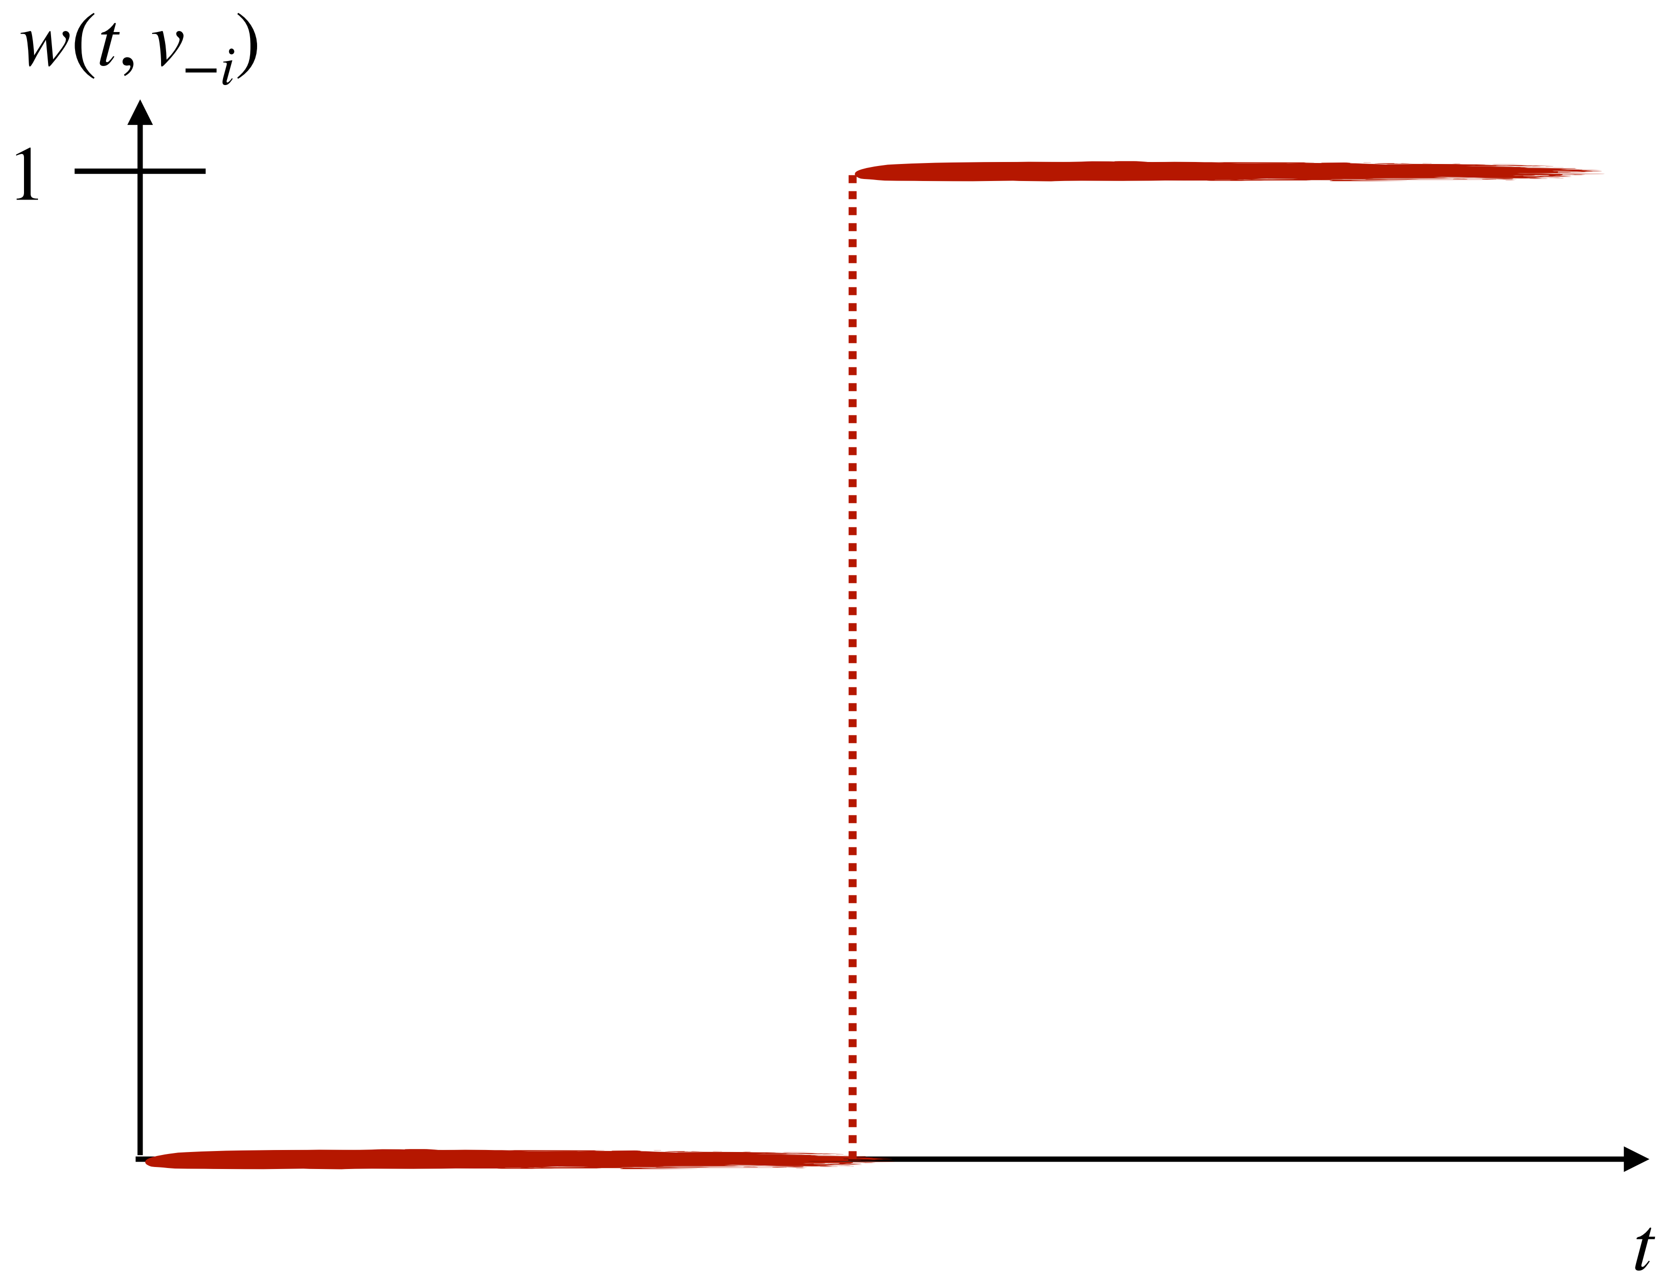

Some pictures

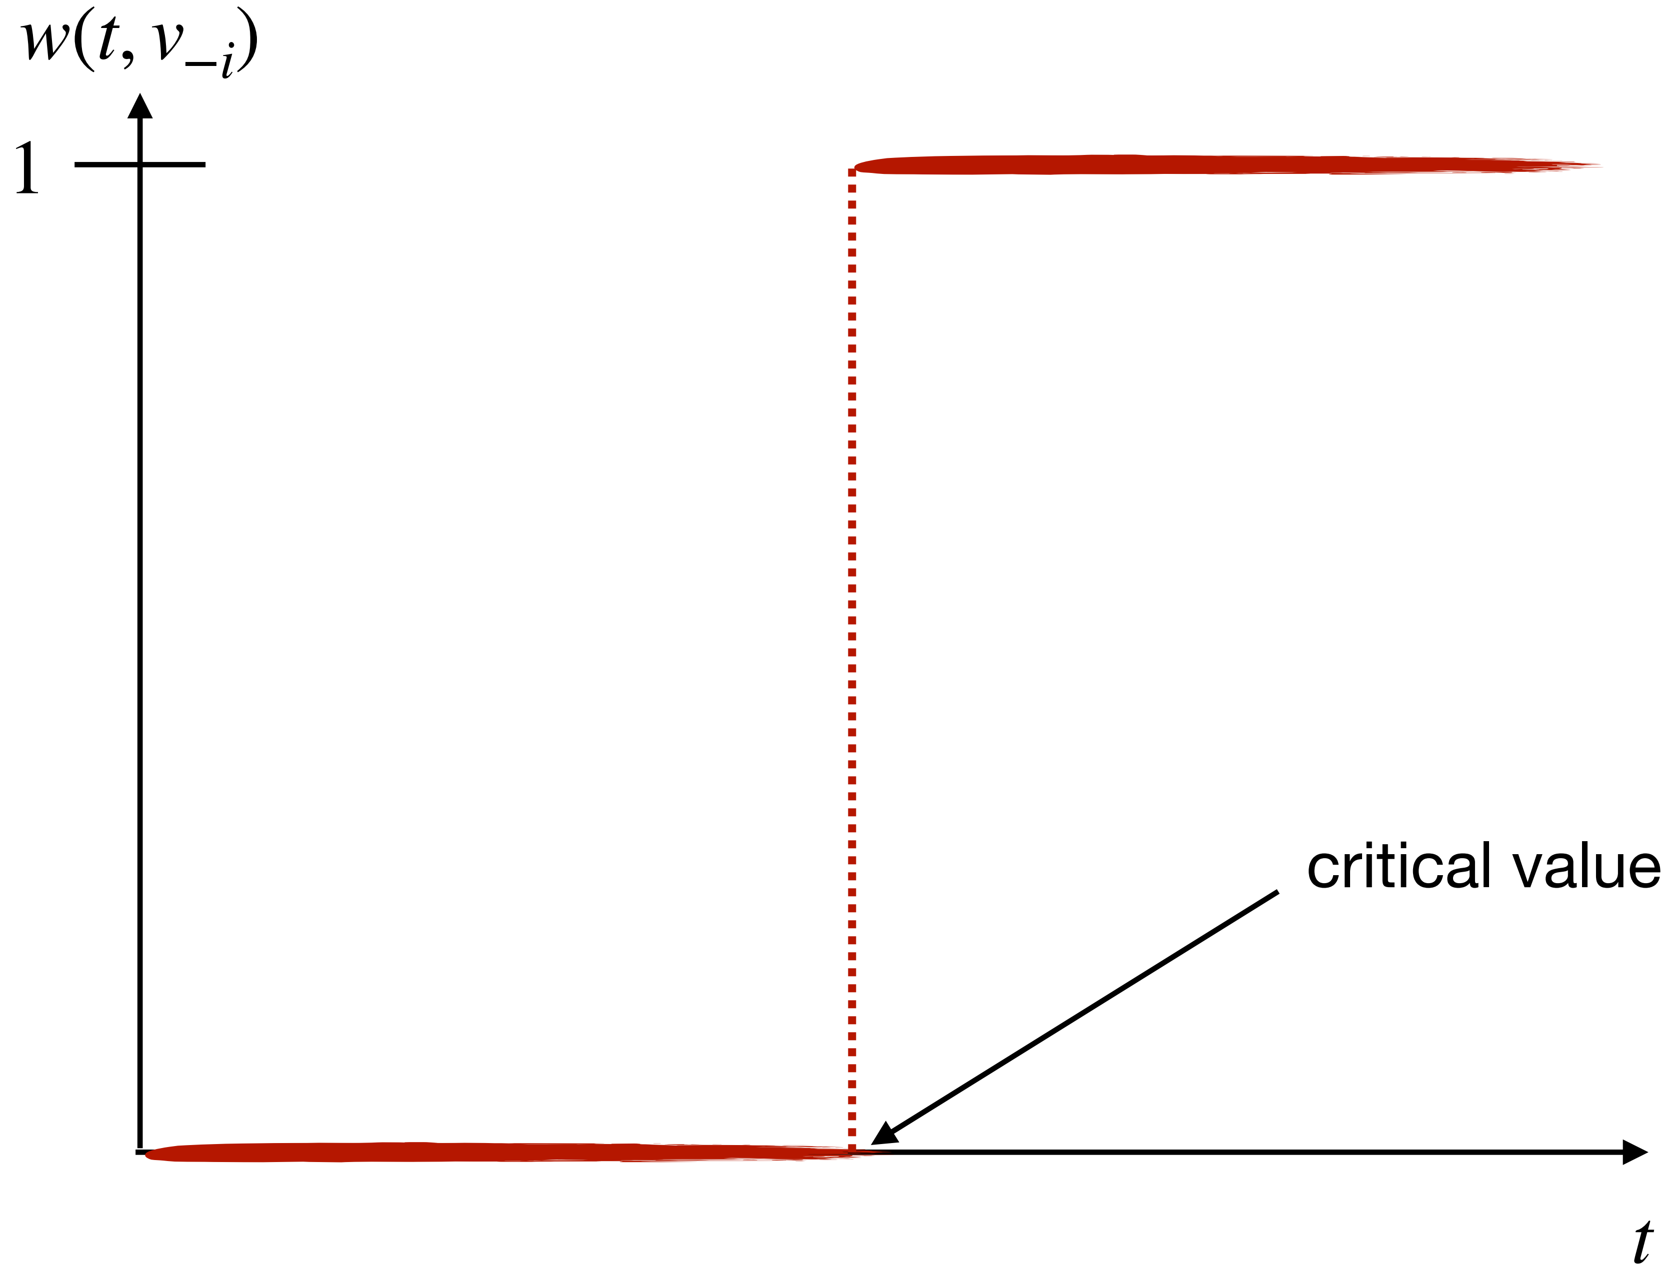

Some pictures

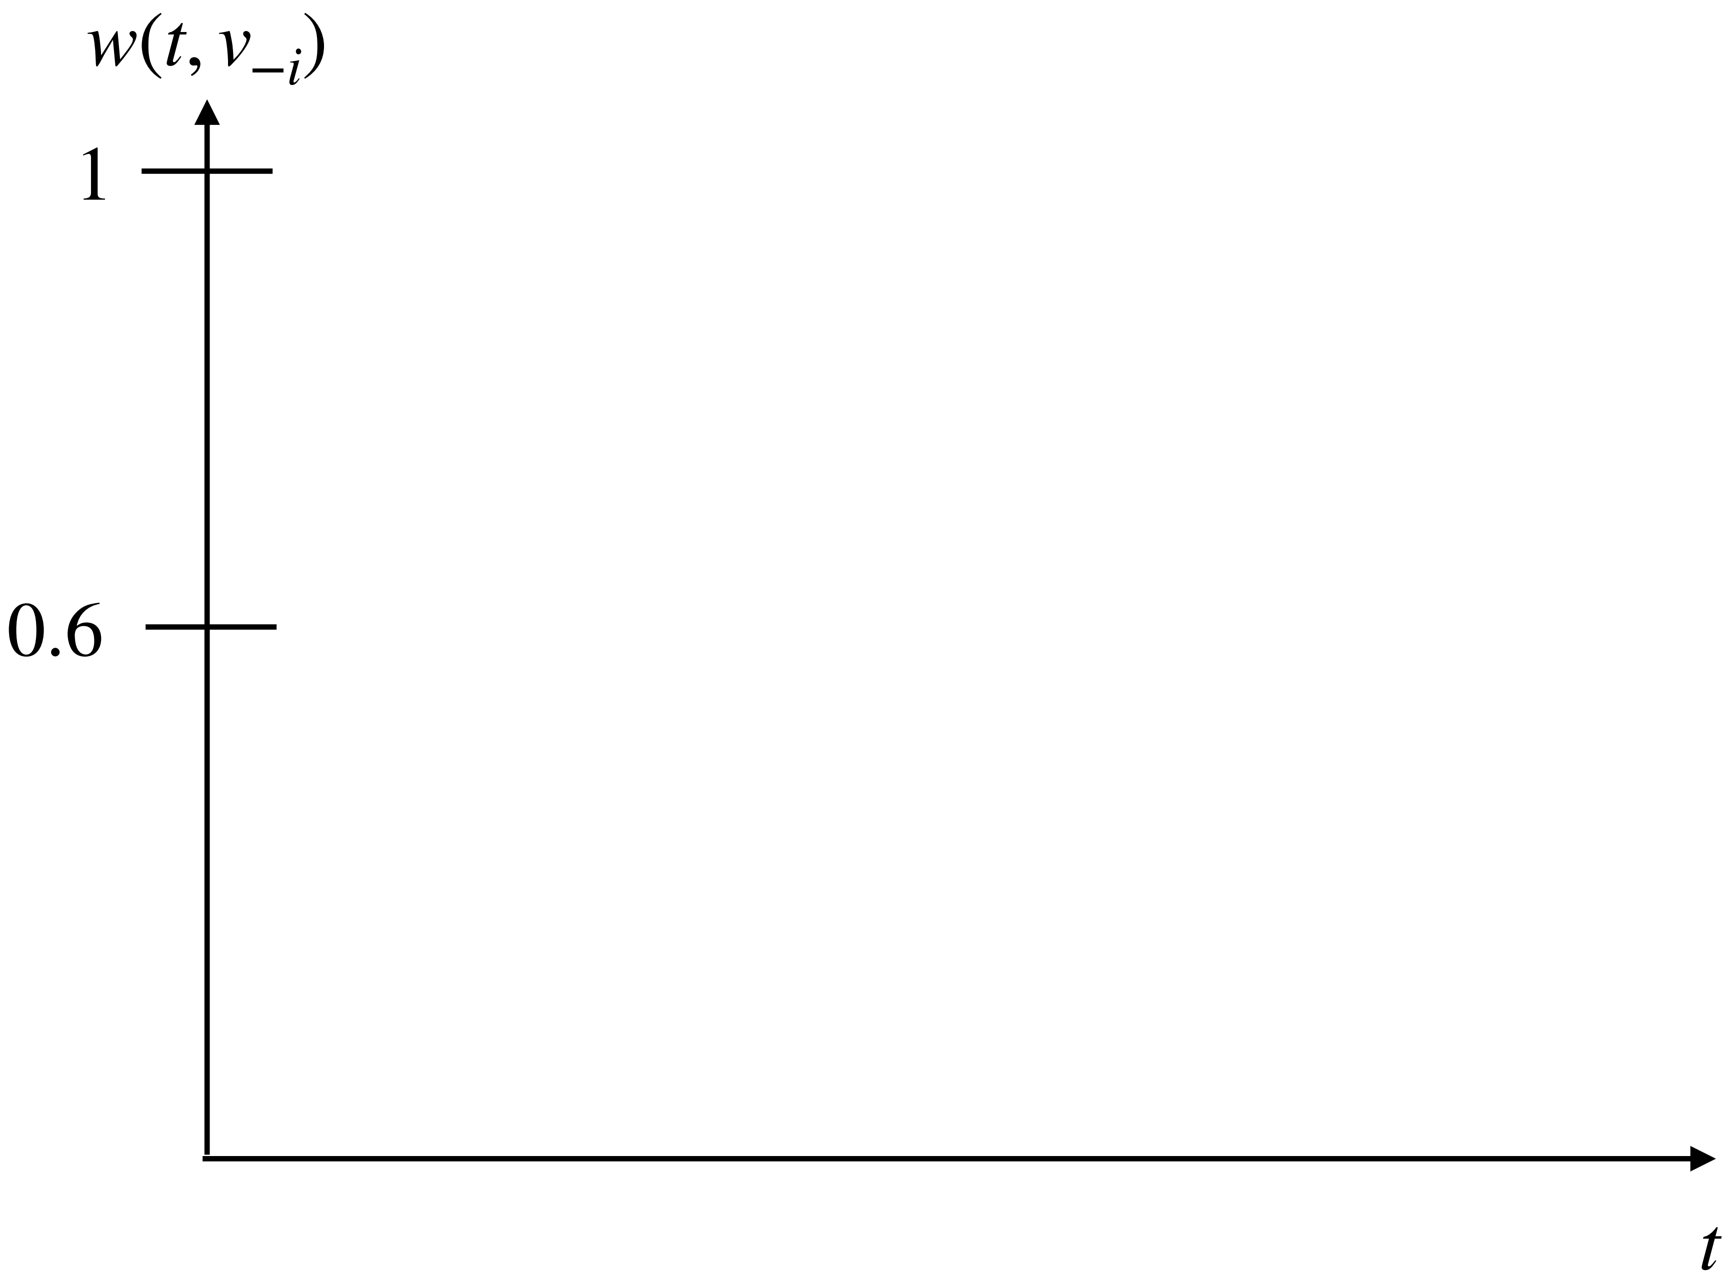

Some pictures

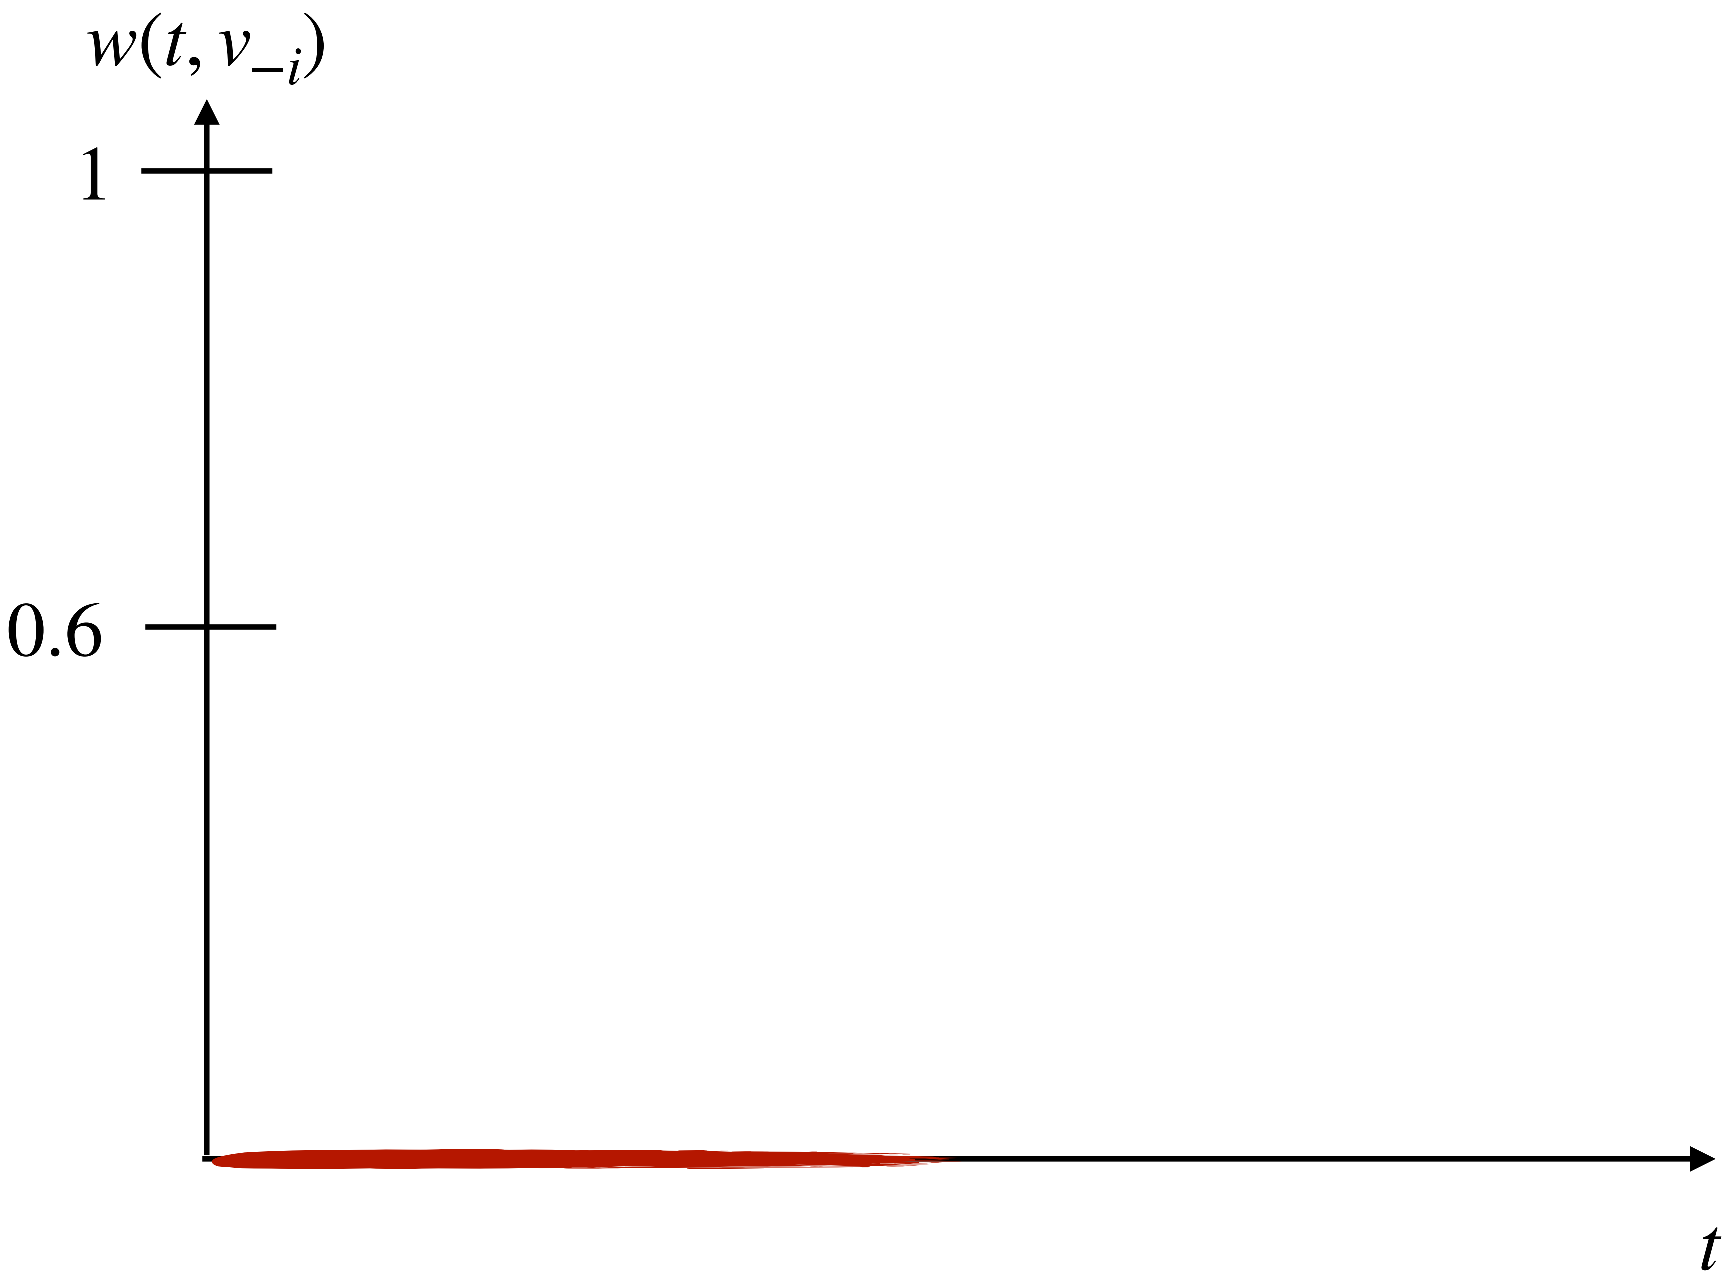

Some pictures

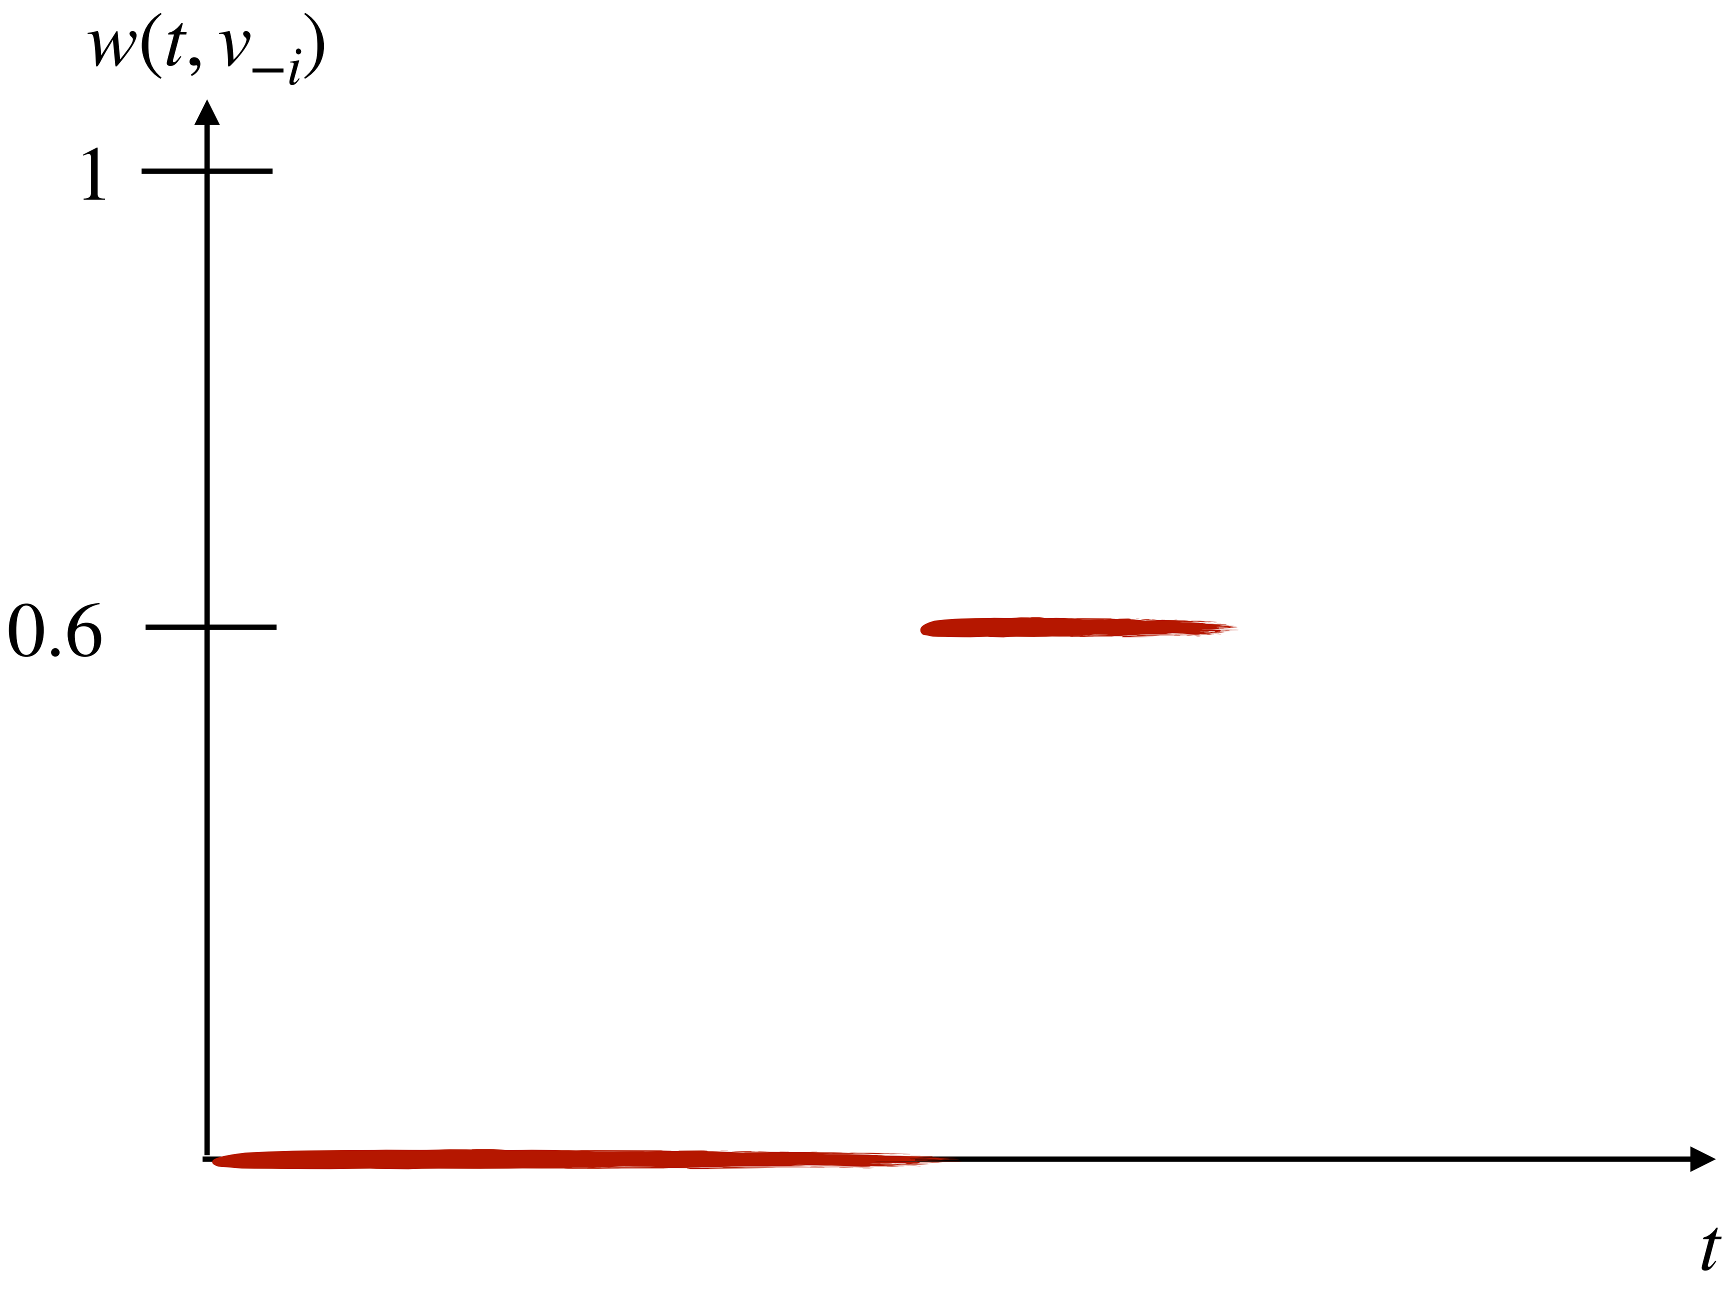

Some pictures

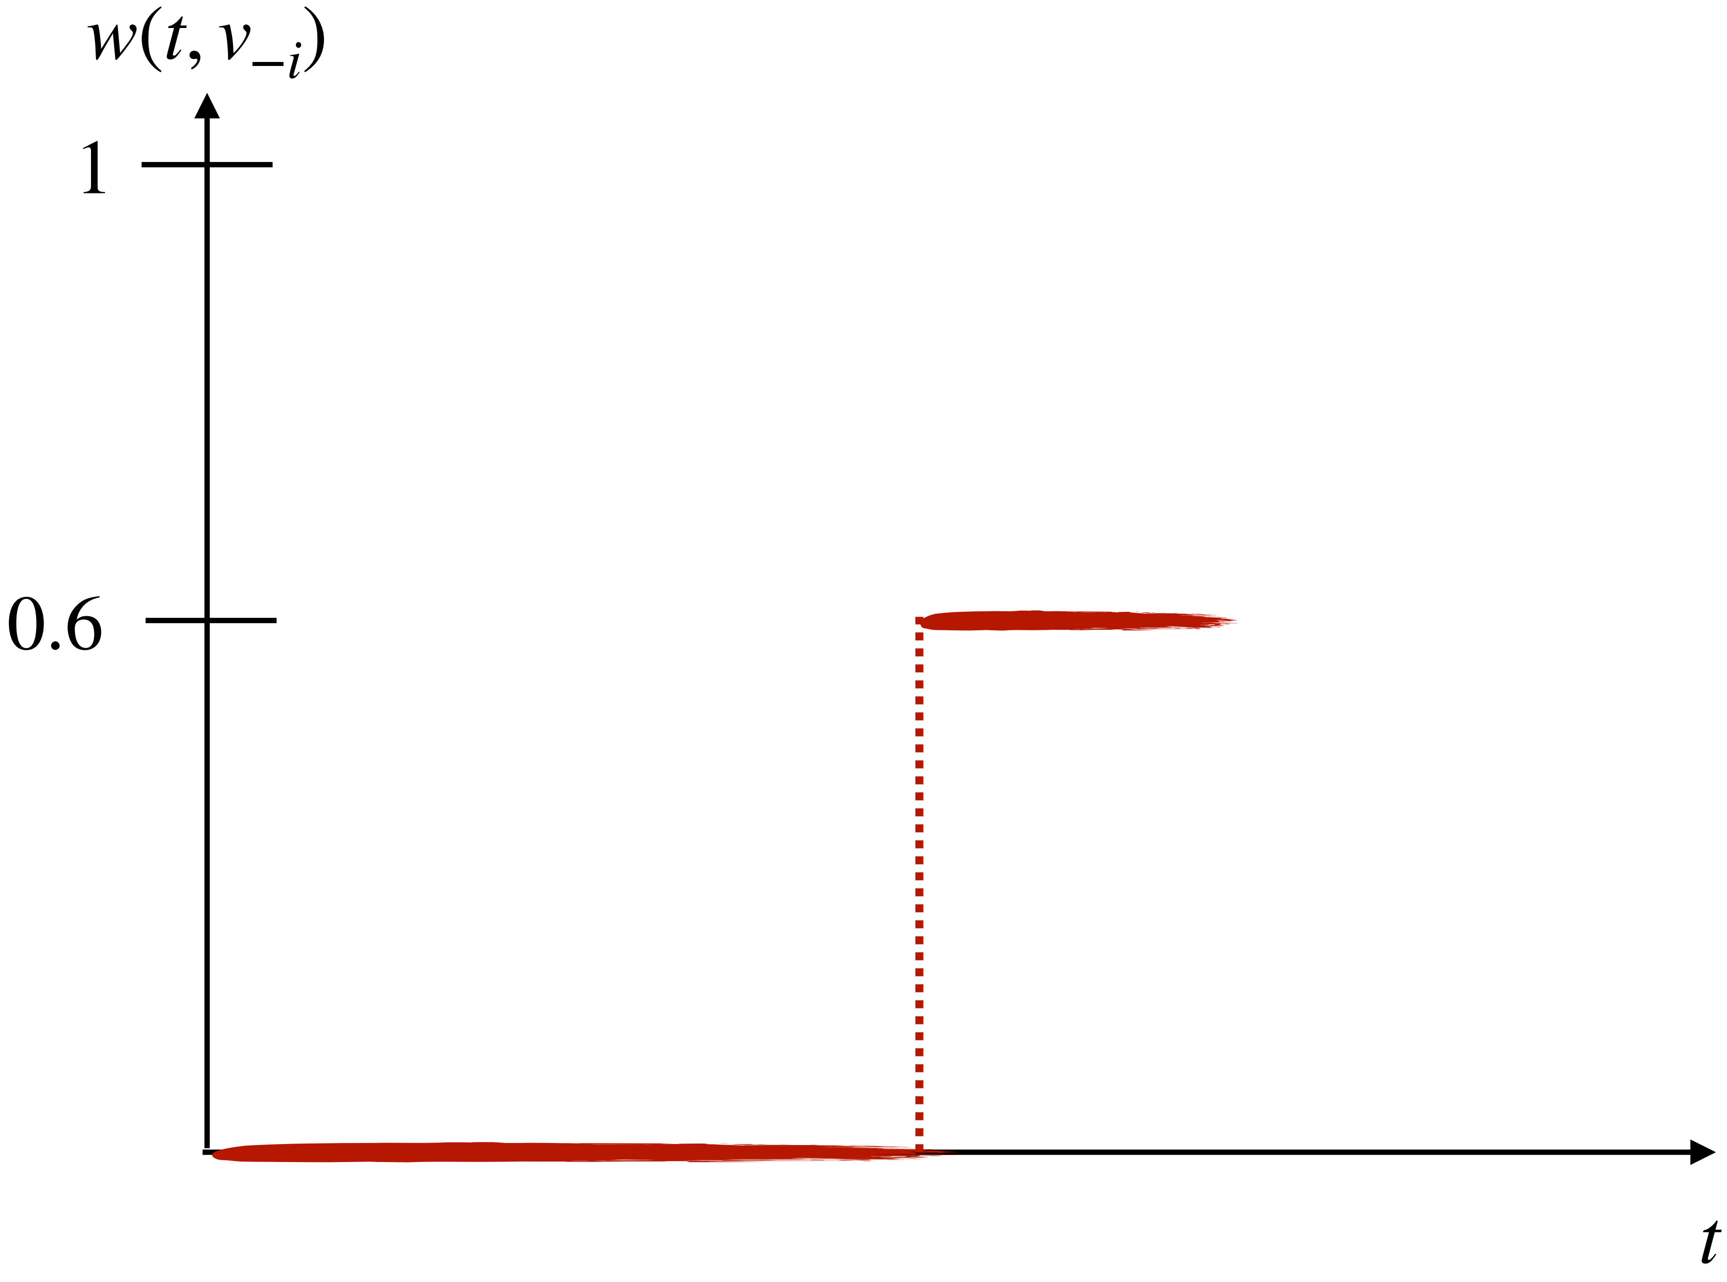

Some pictures

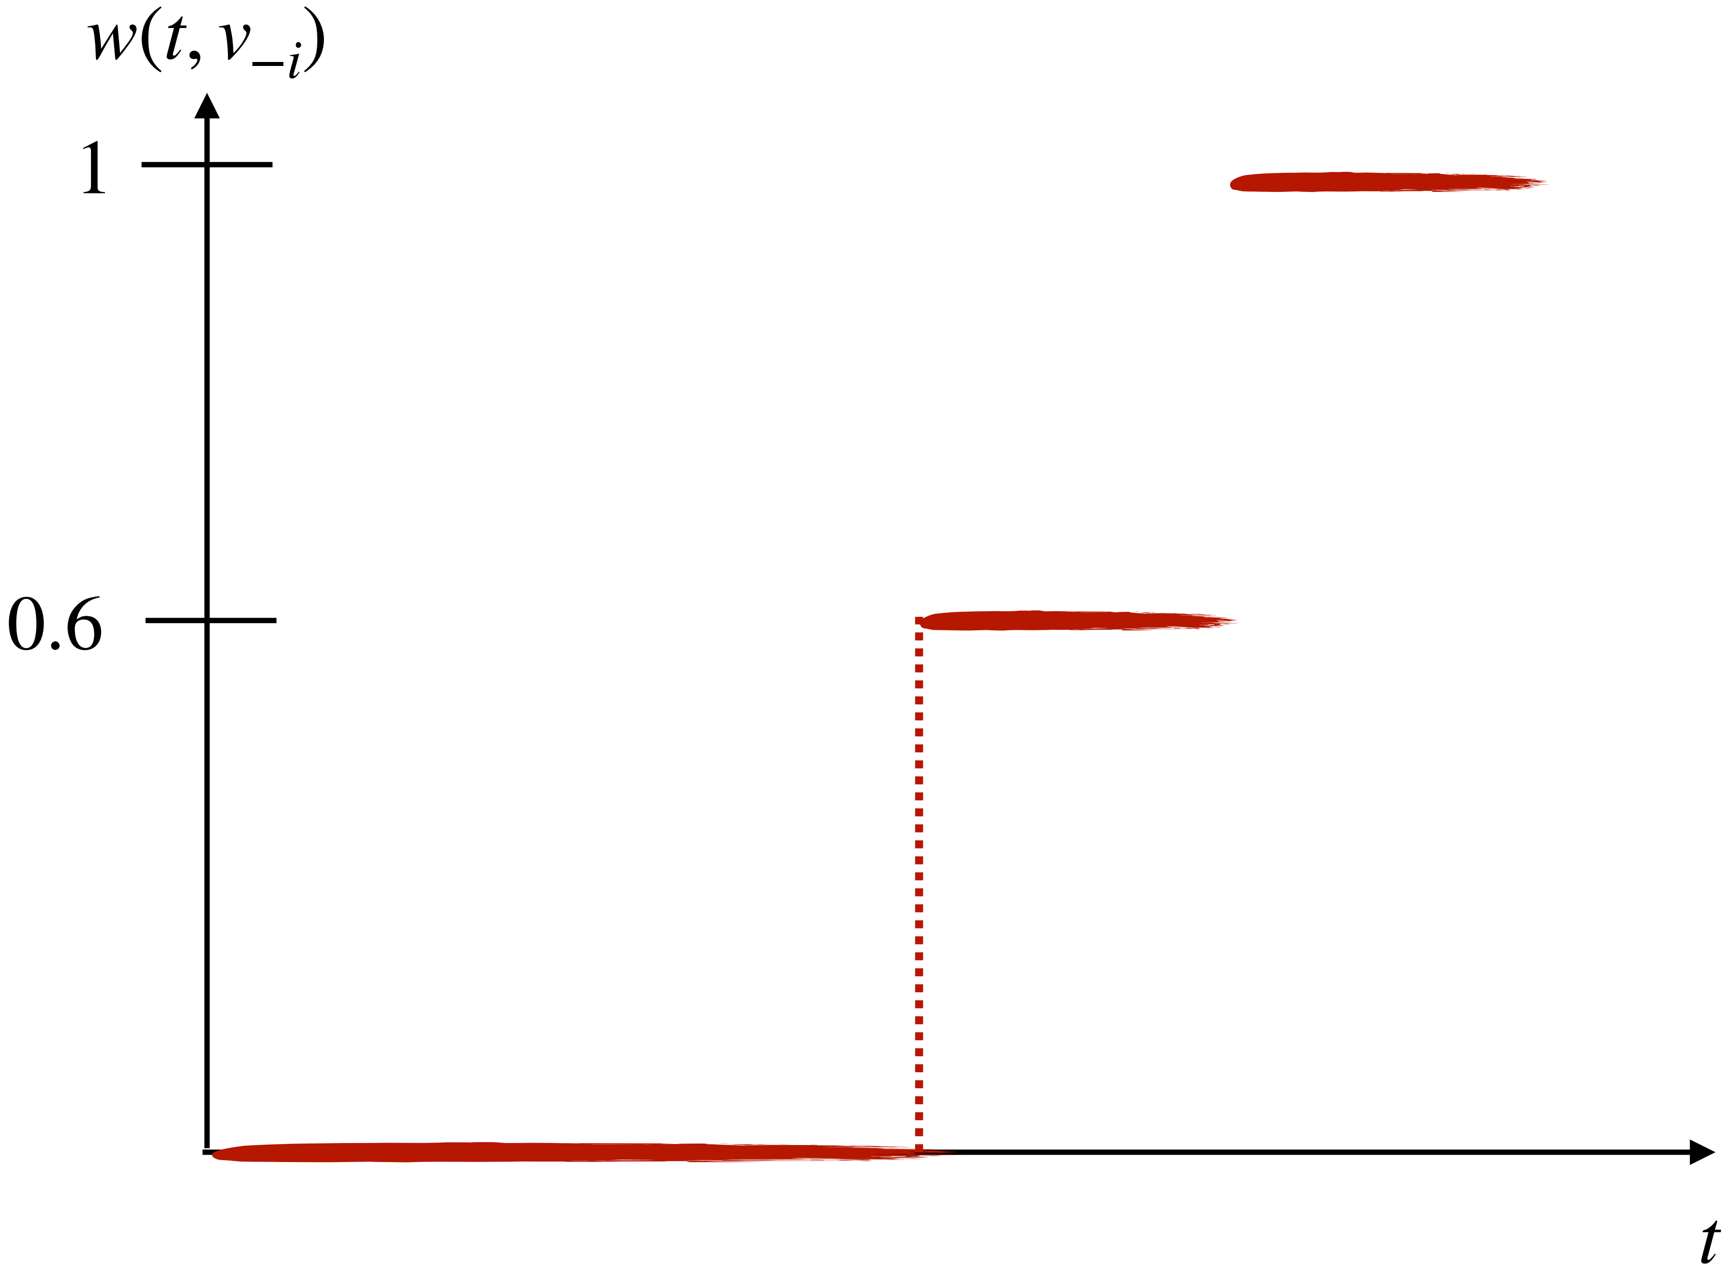

Some pictures

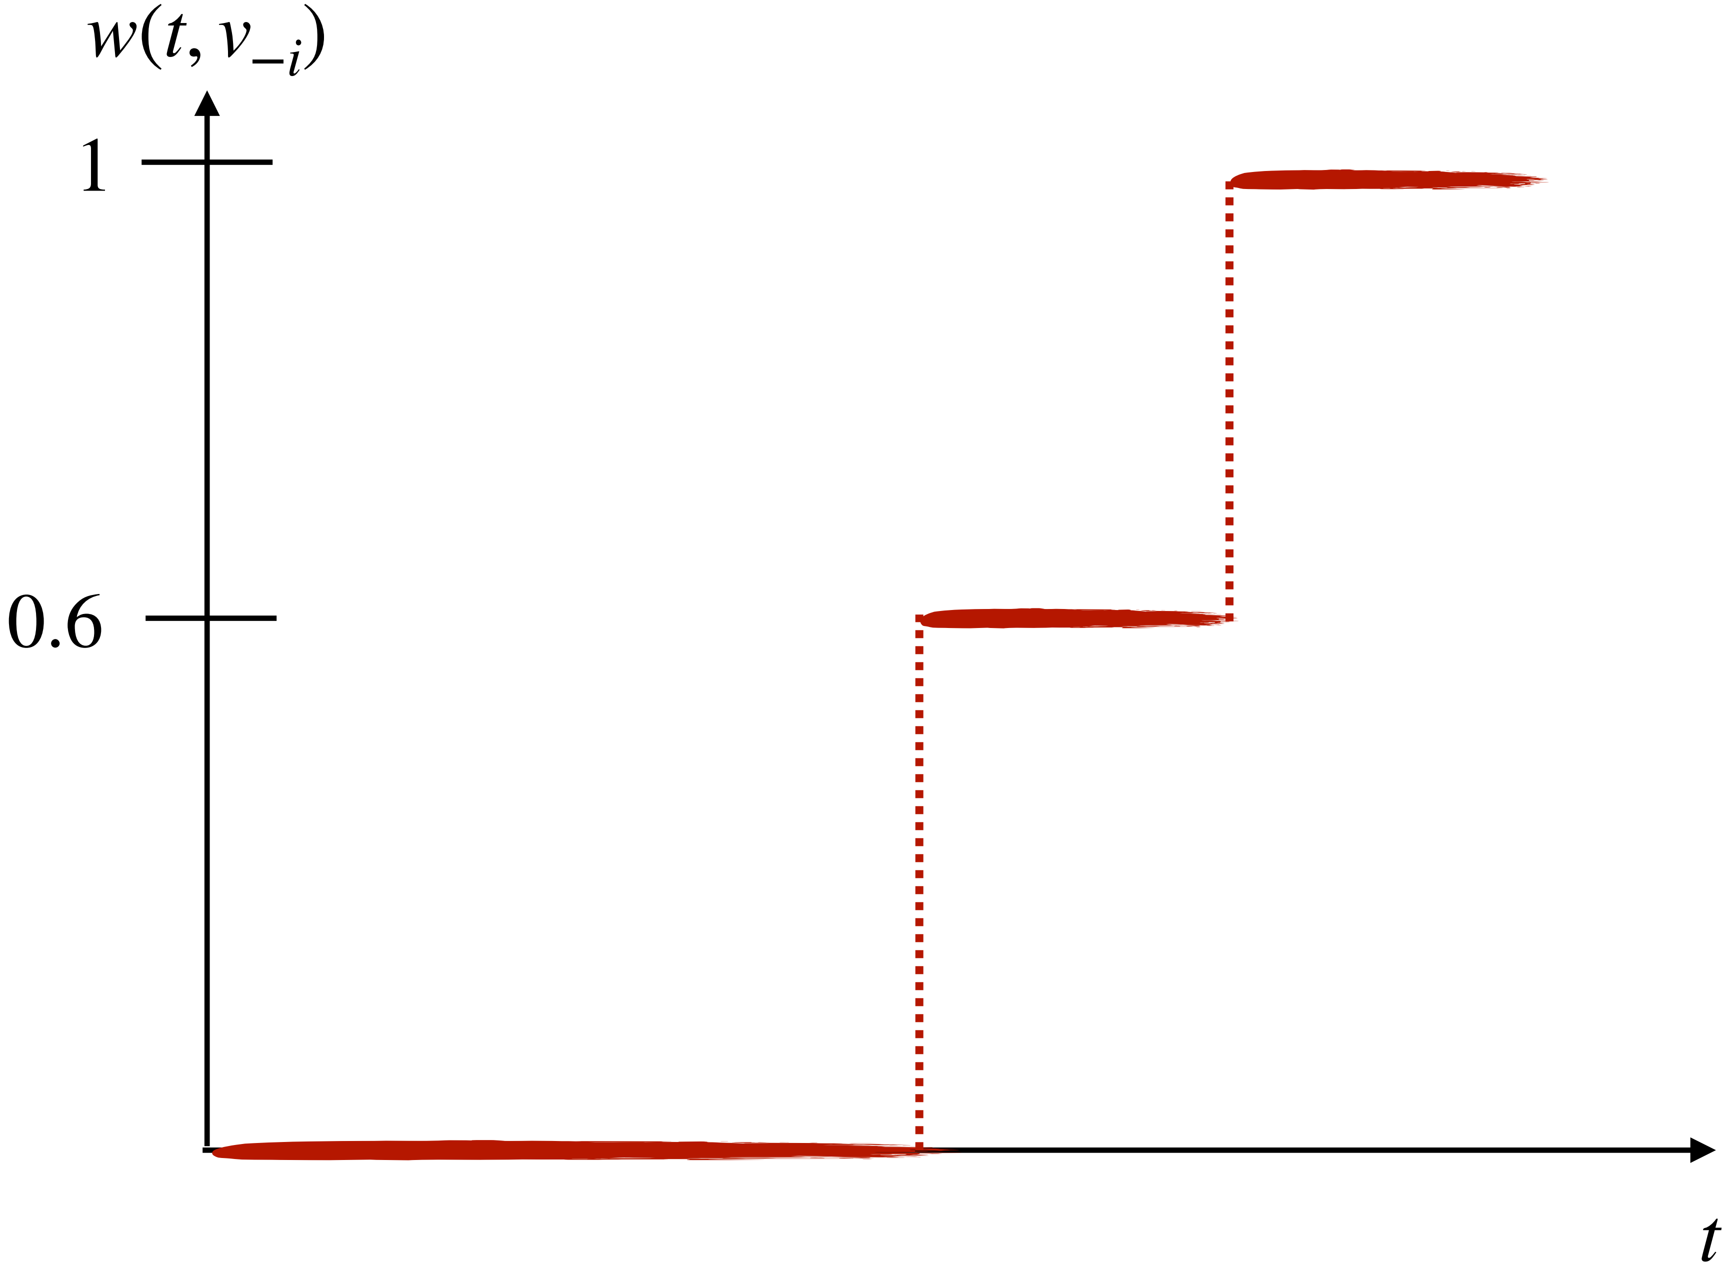

Some pictures

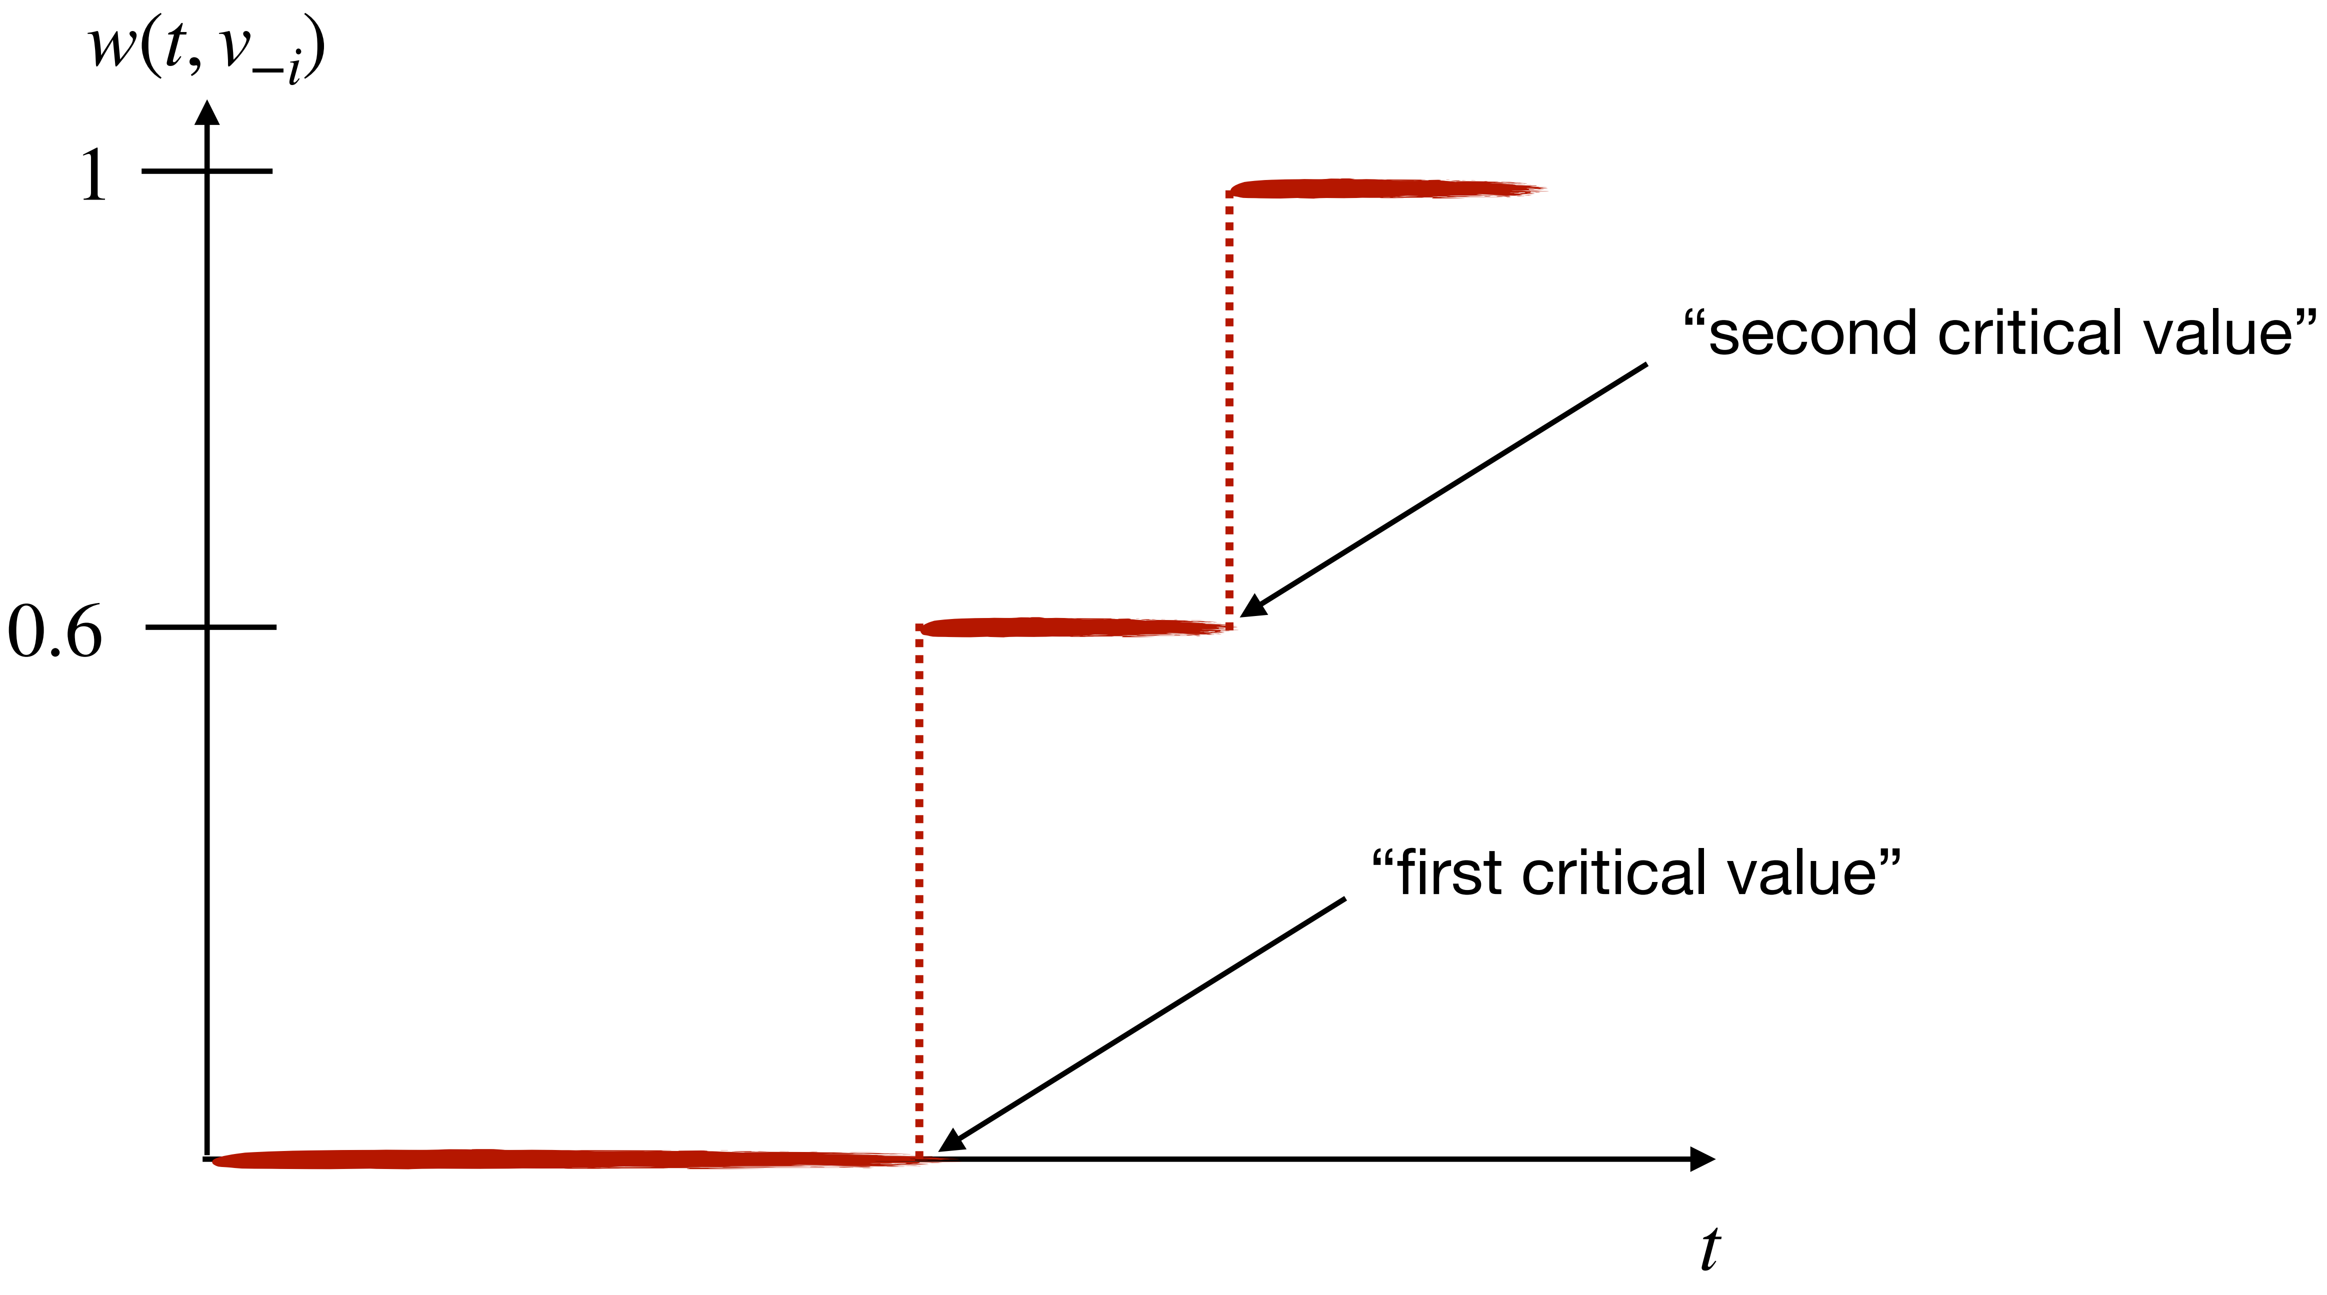

Some pictures

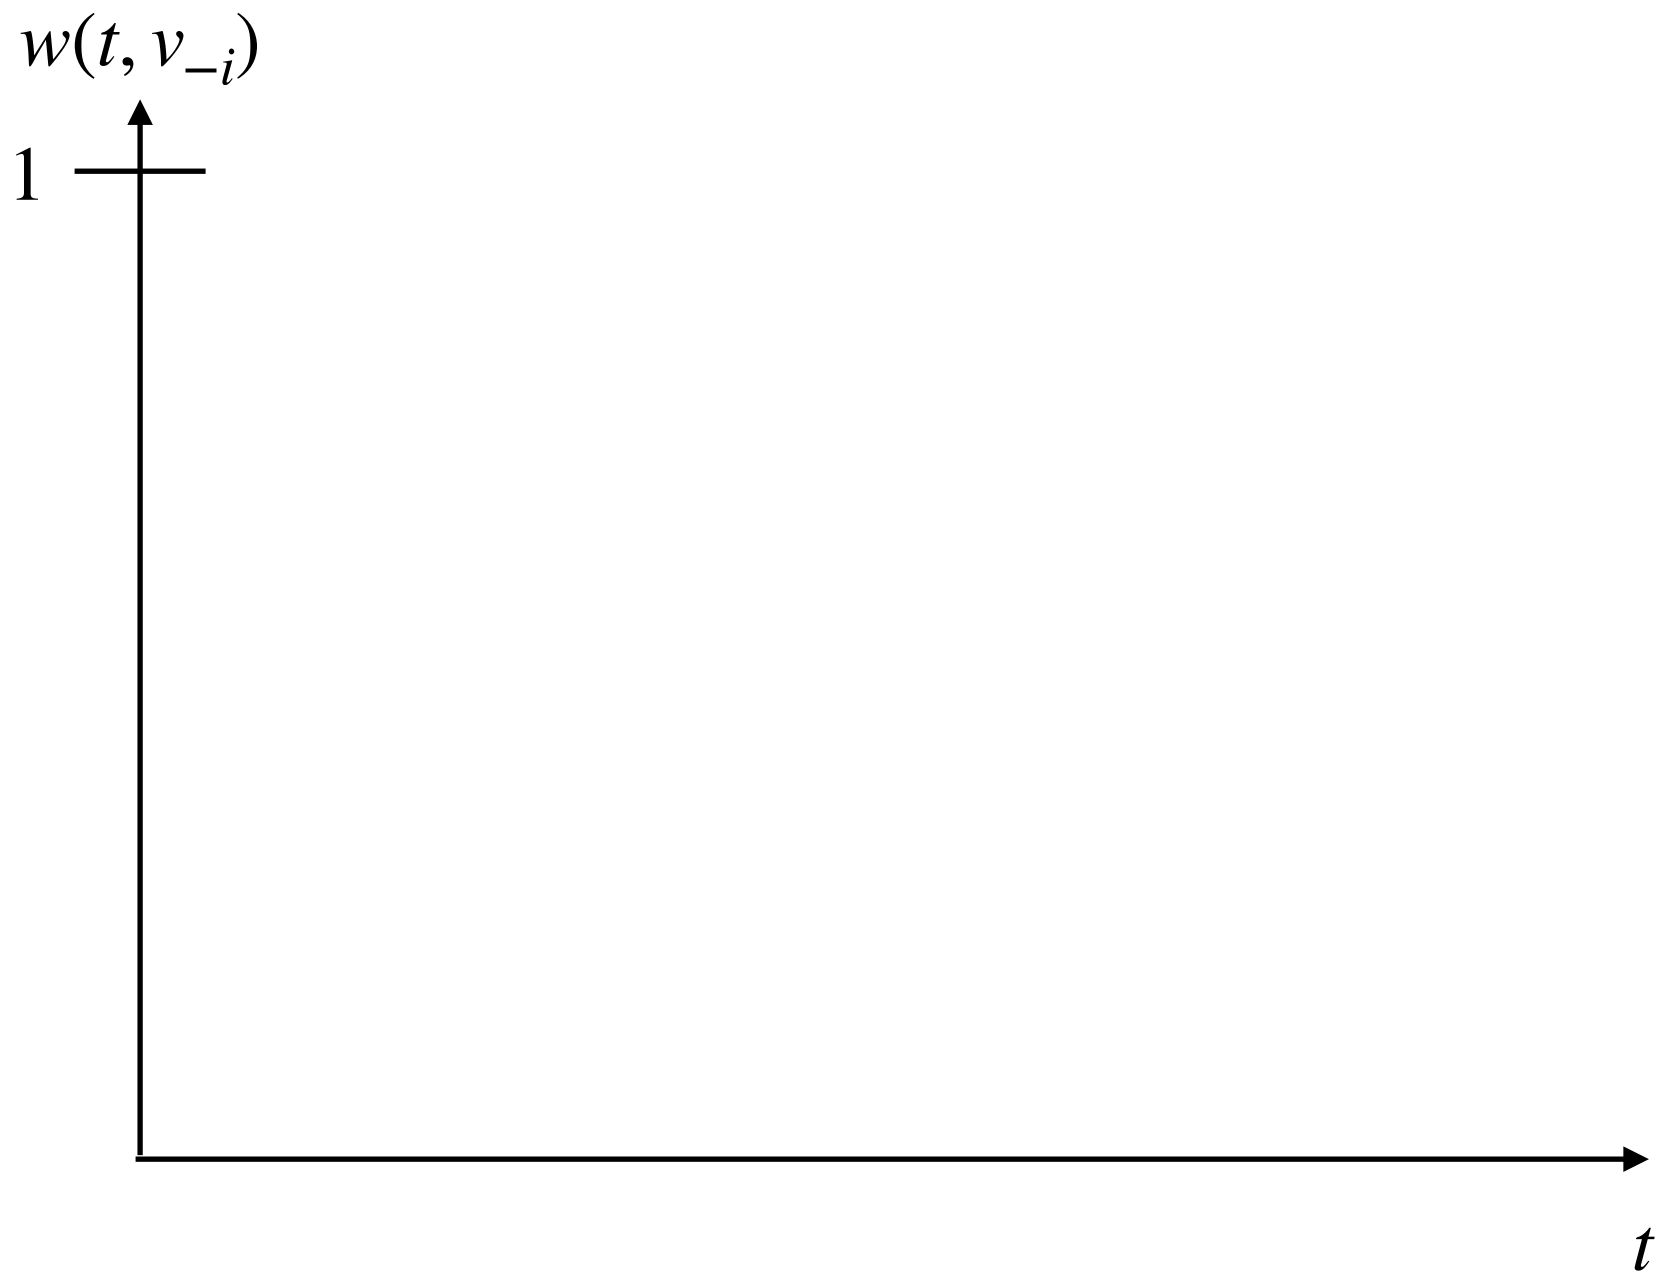

Some pictures

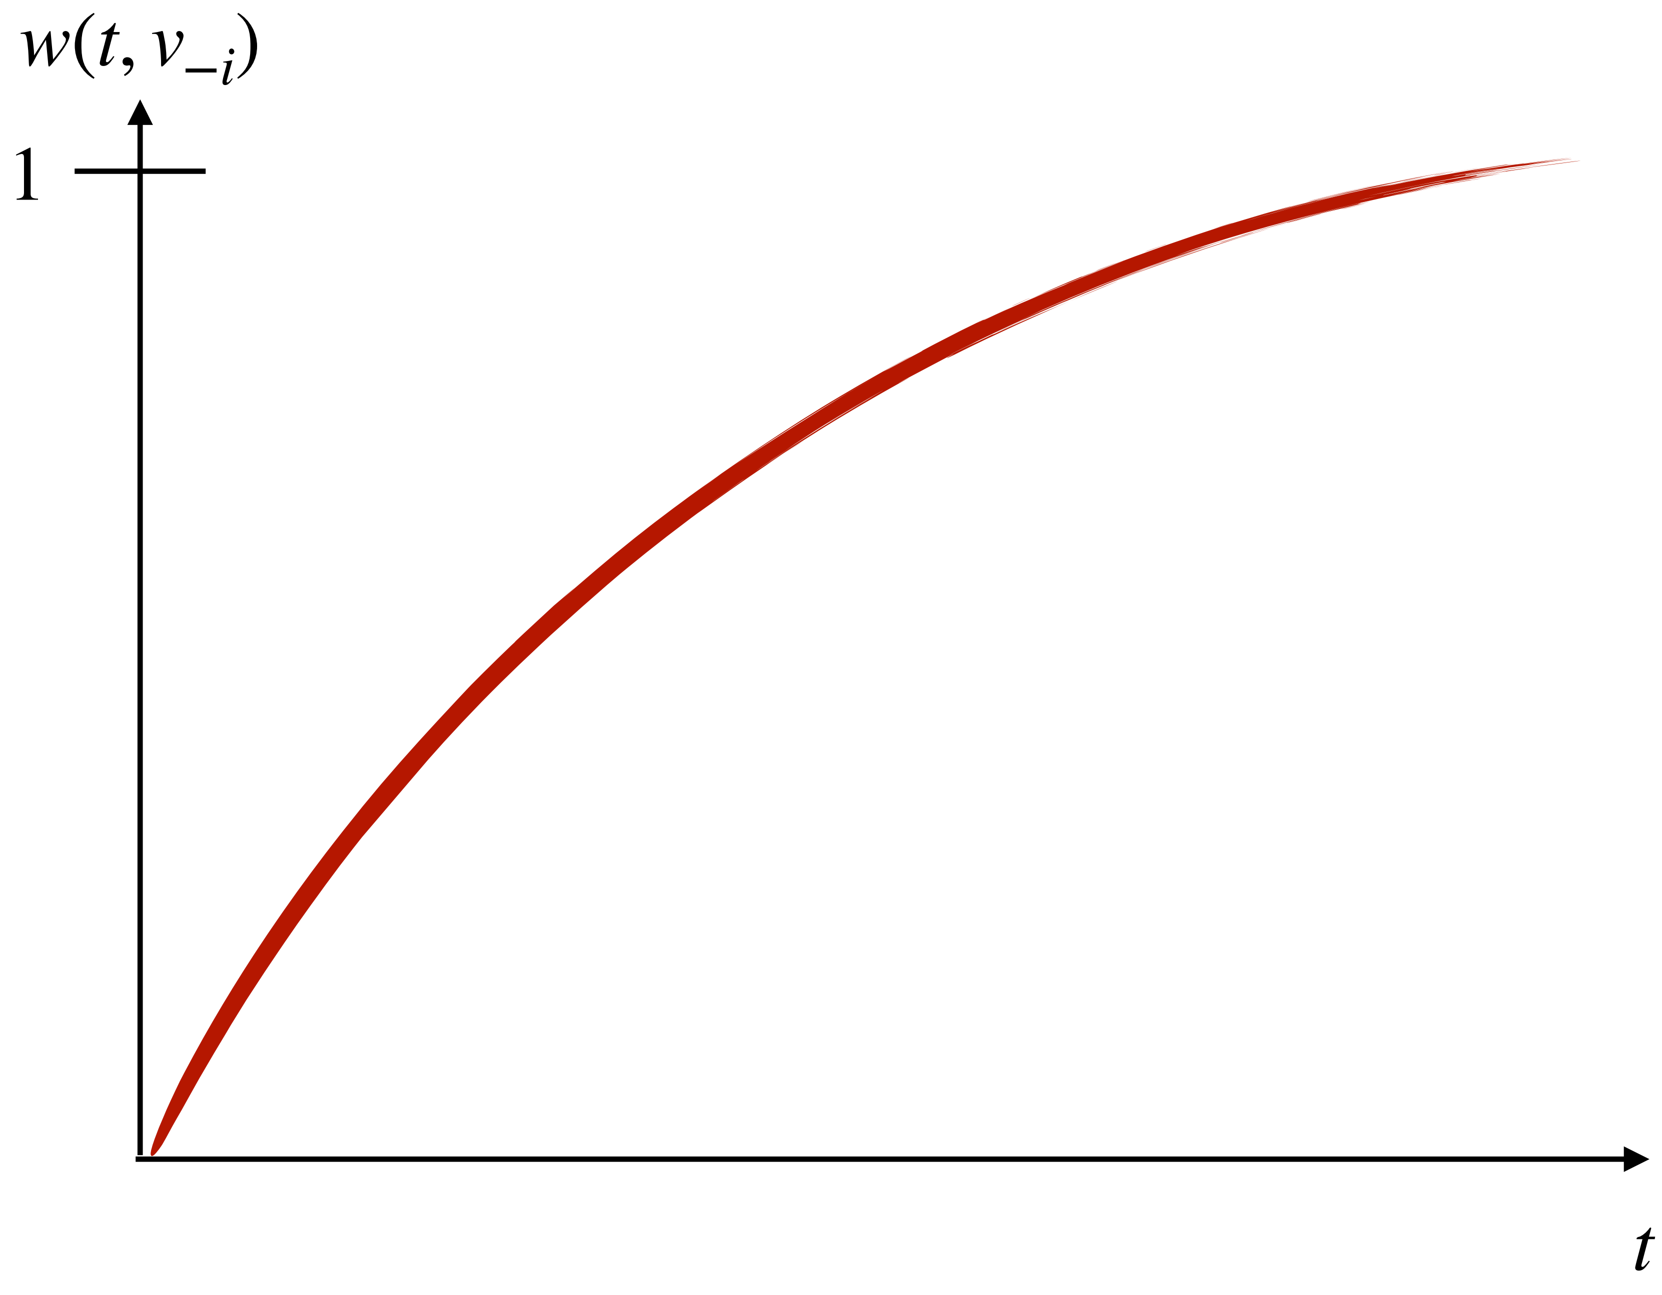

Some pictures

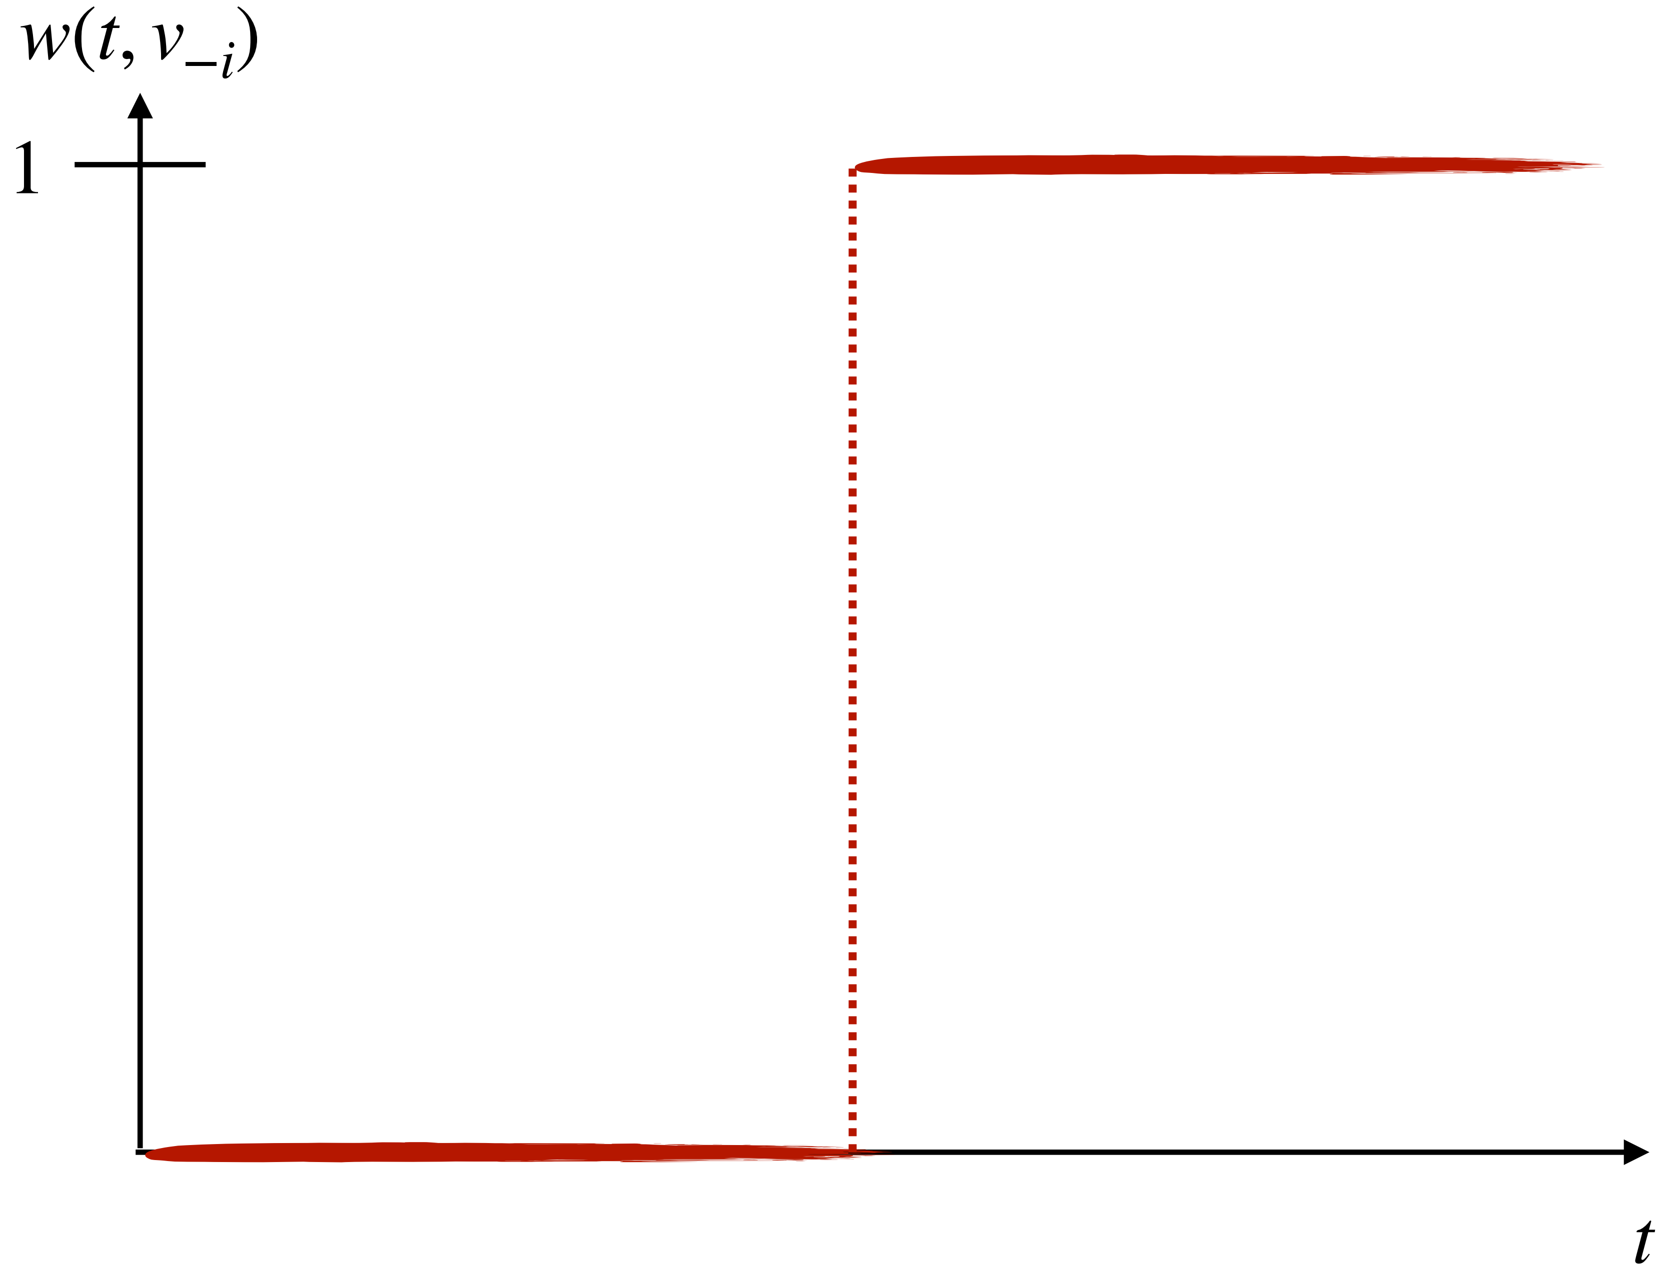

Some pictures

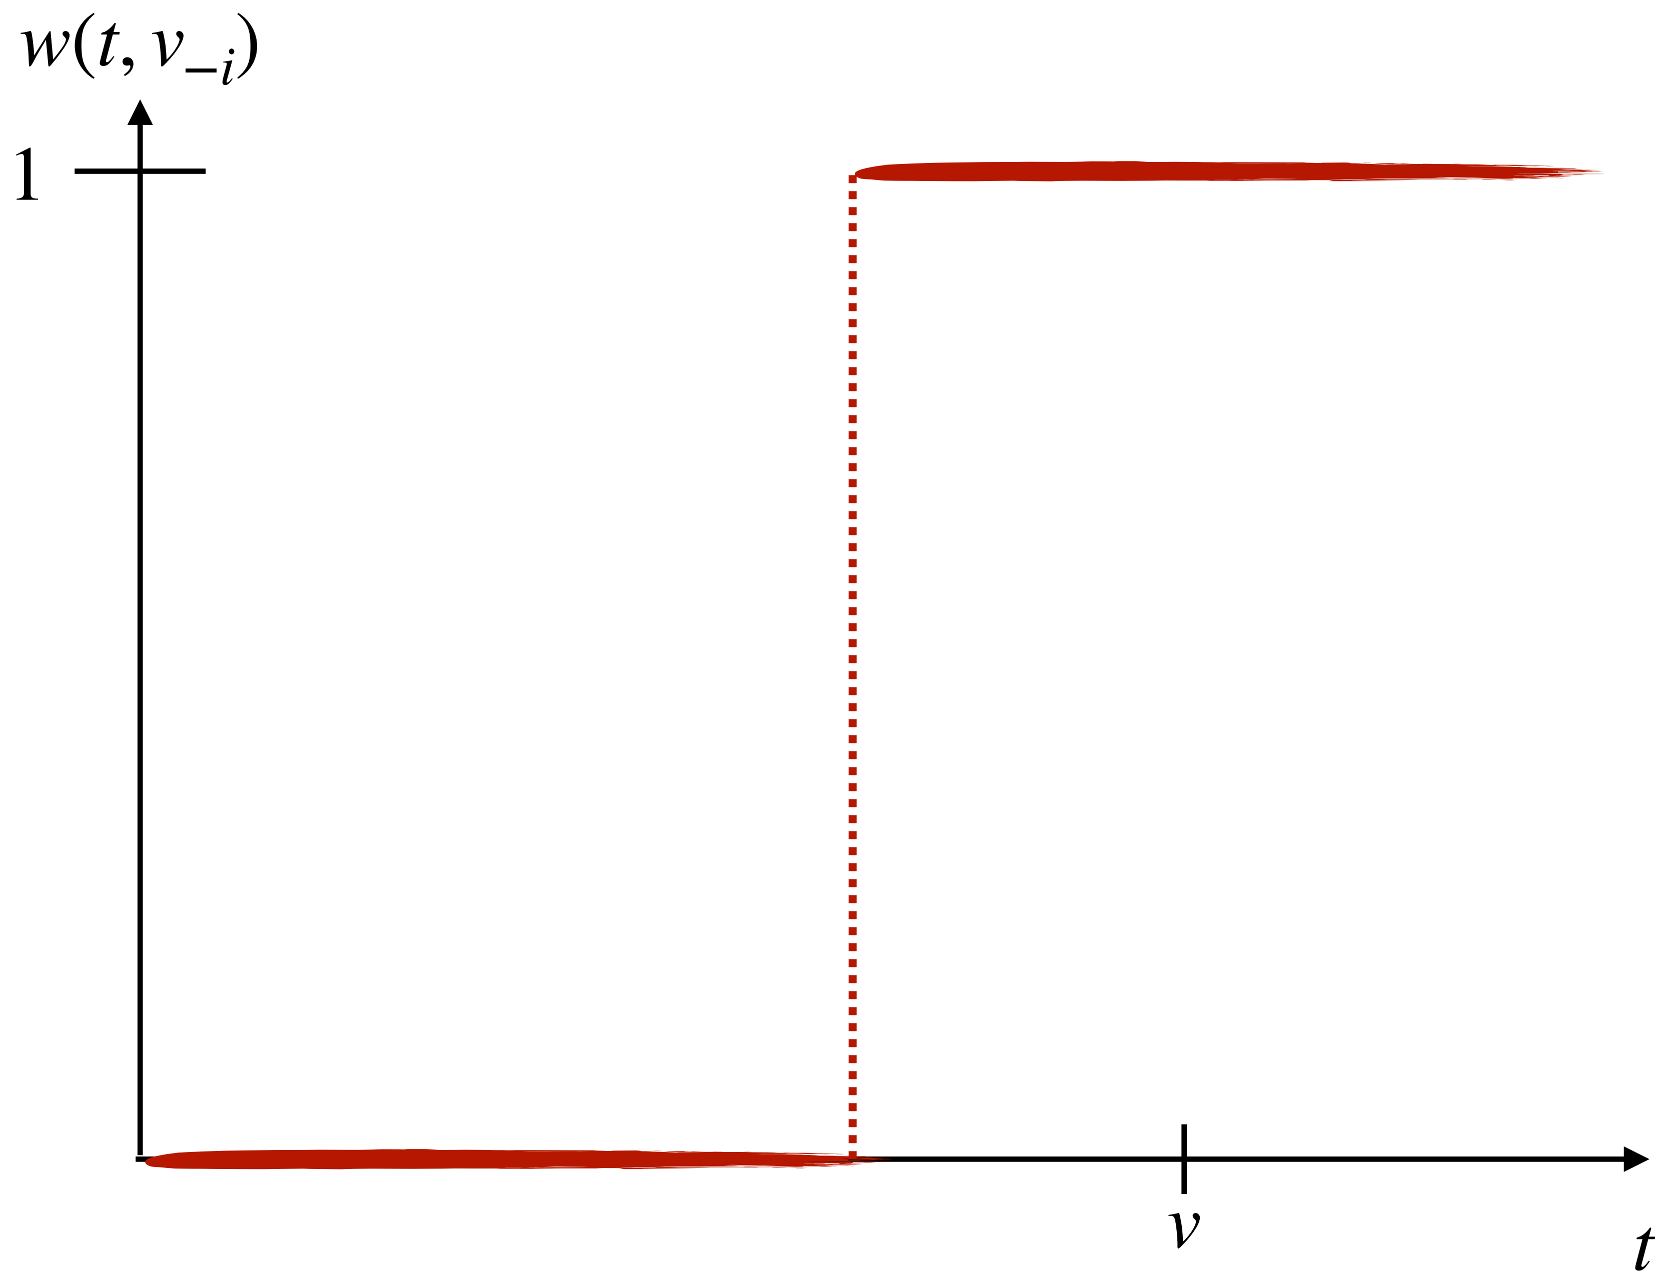

Some pictures

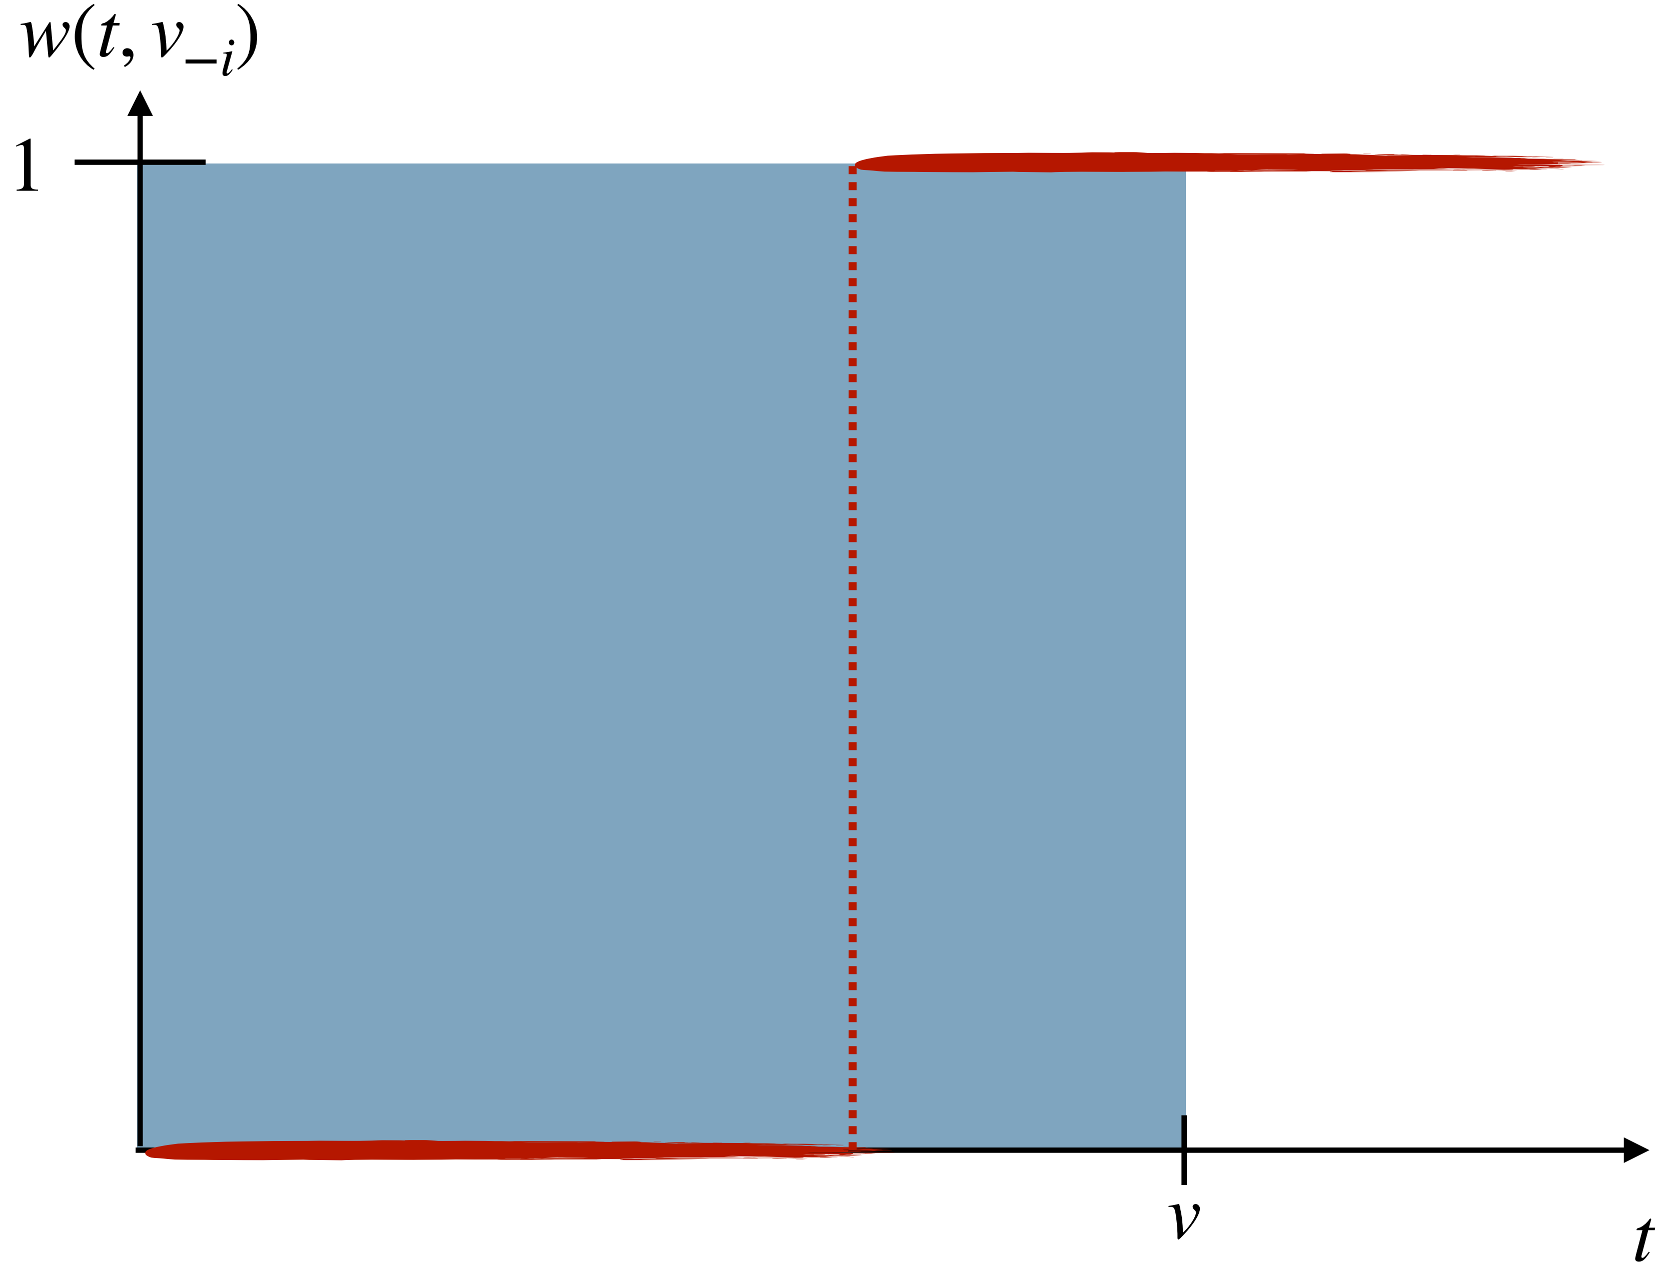

Some pictures

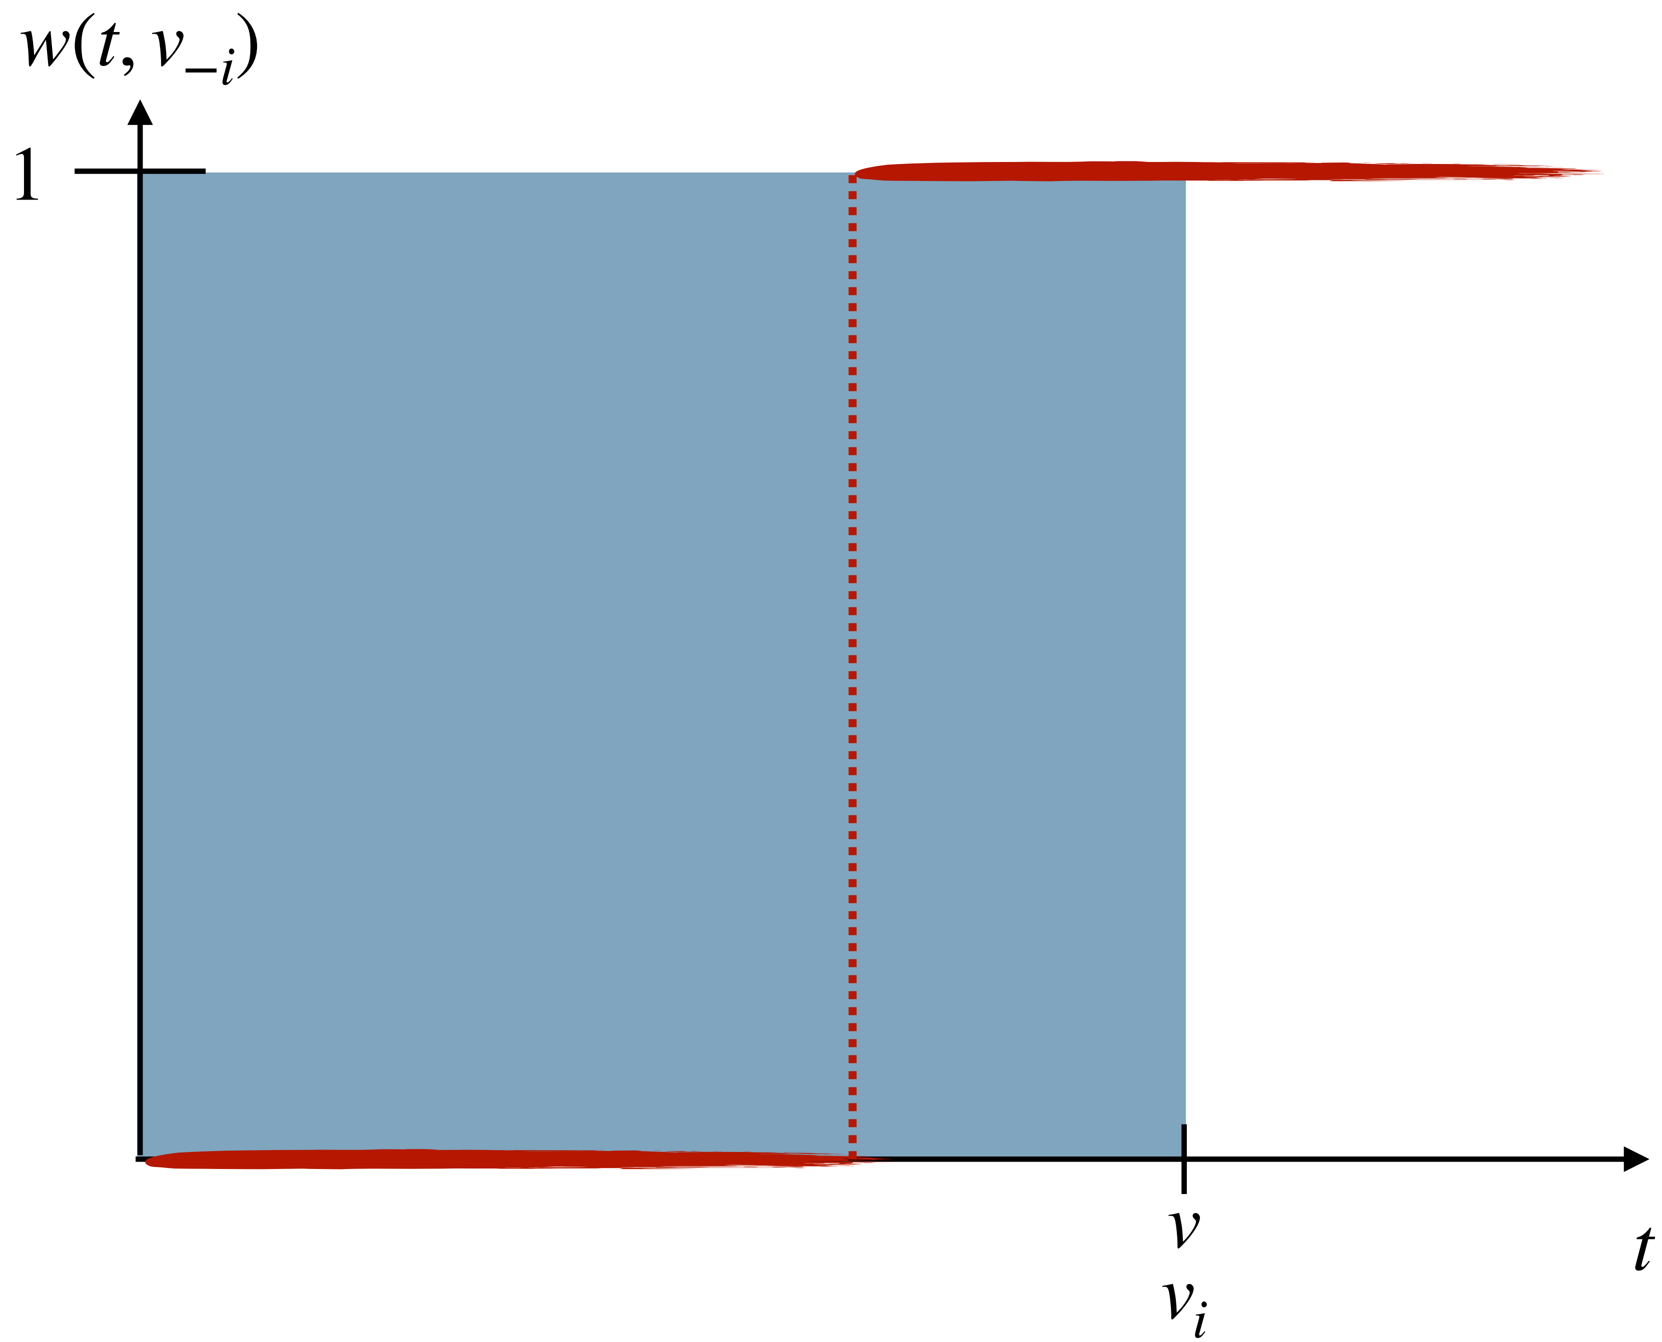

Some pictures

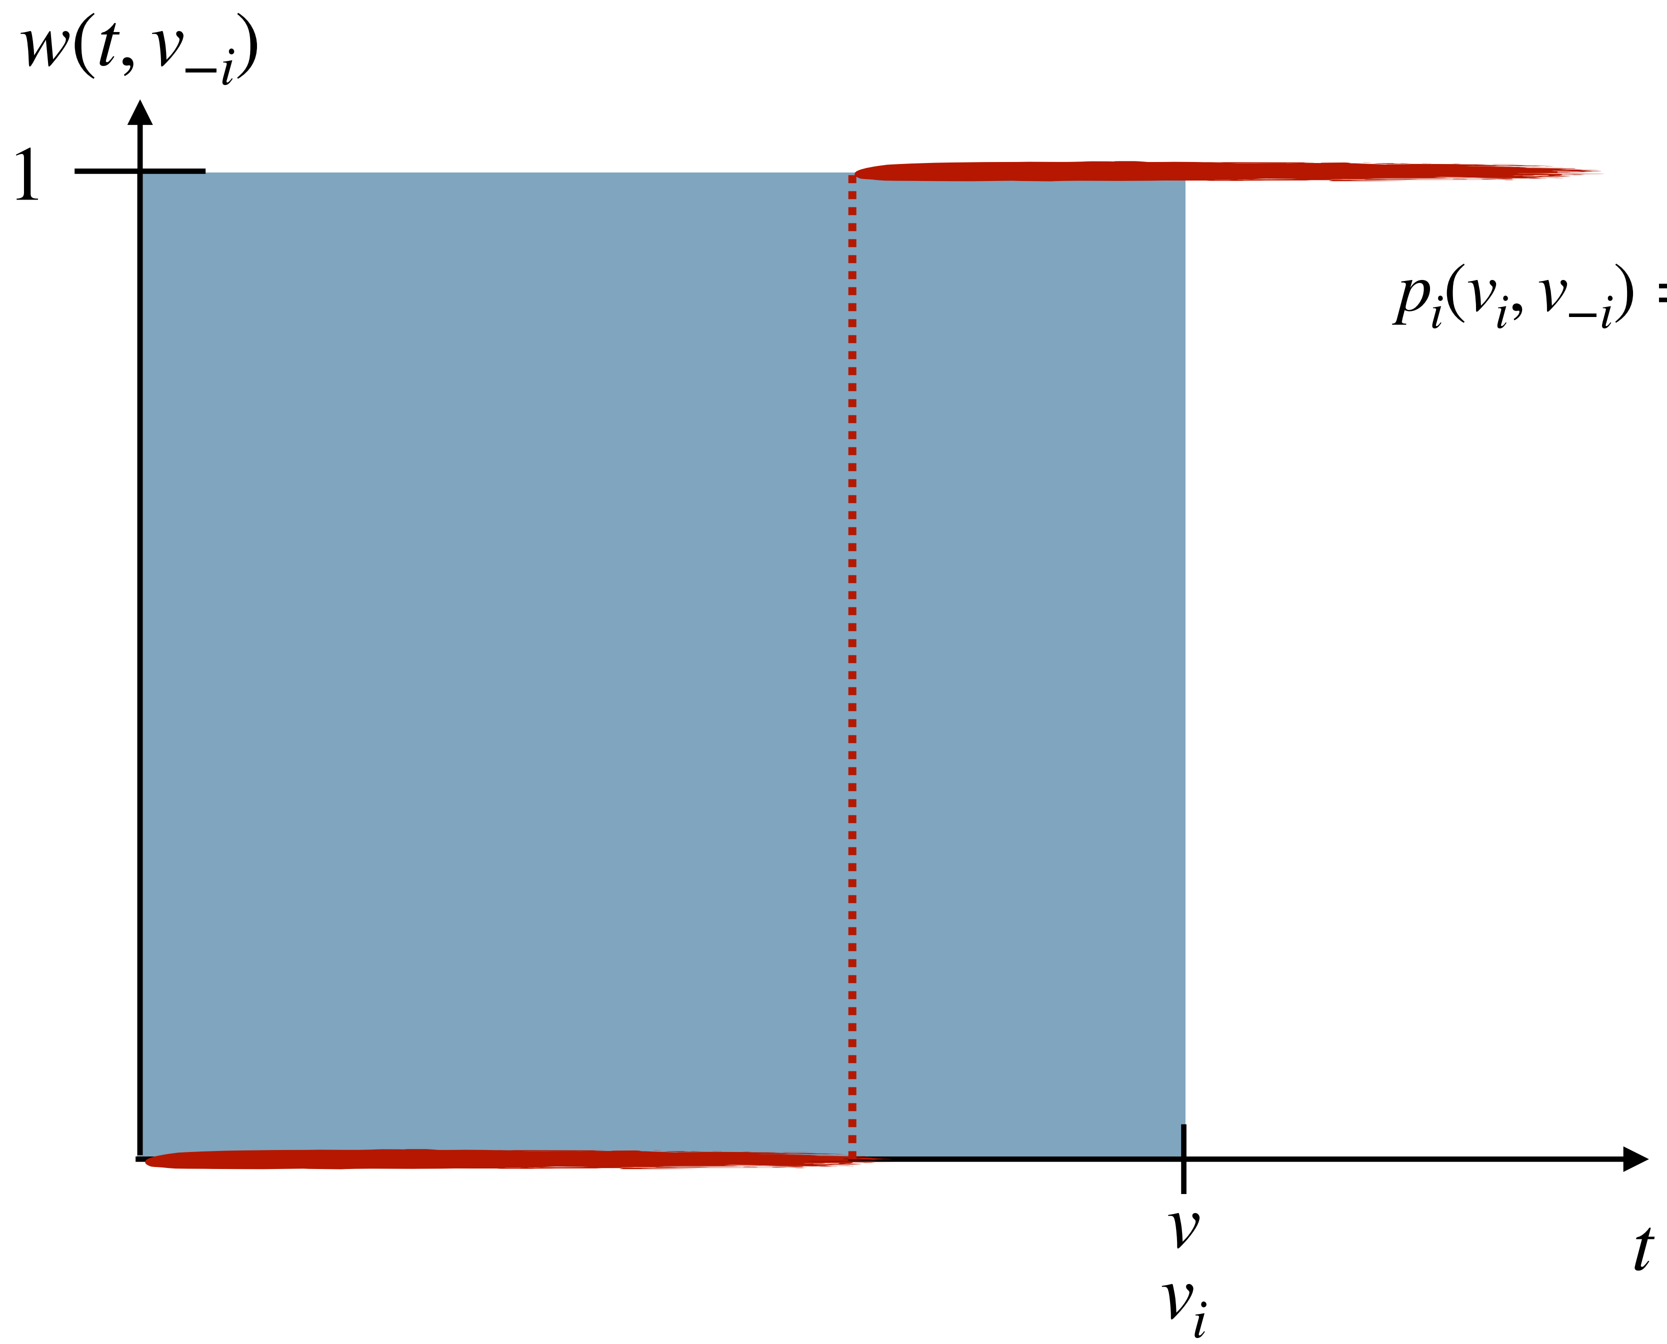

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

Some pictures

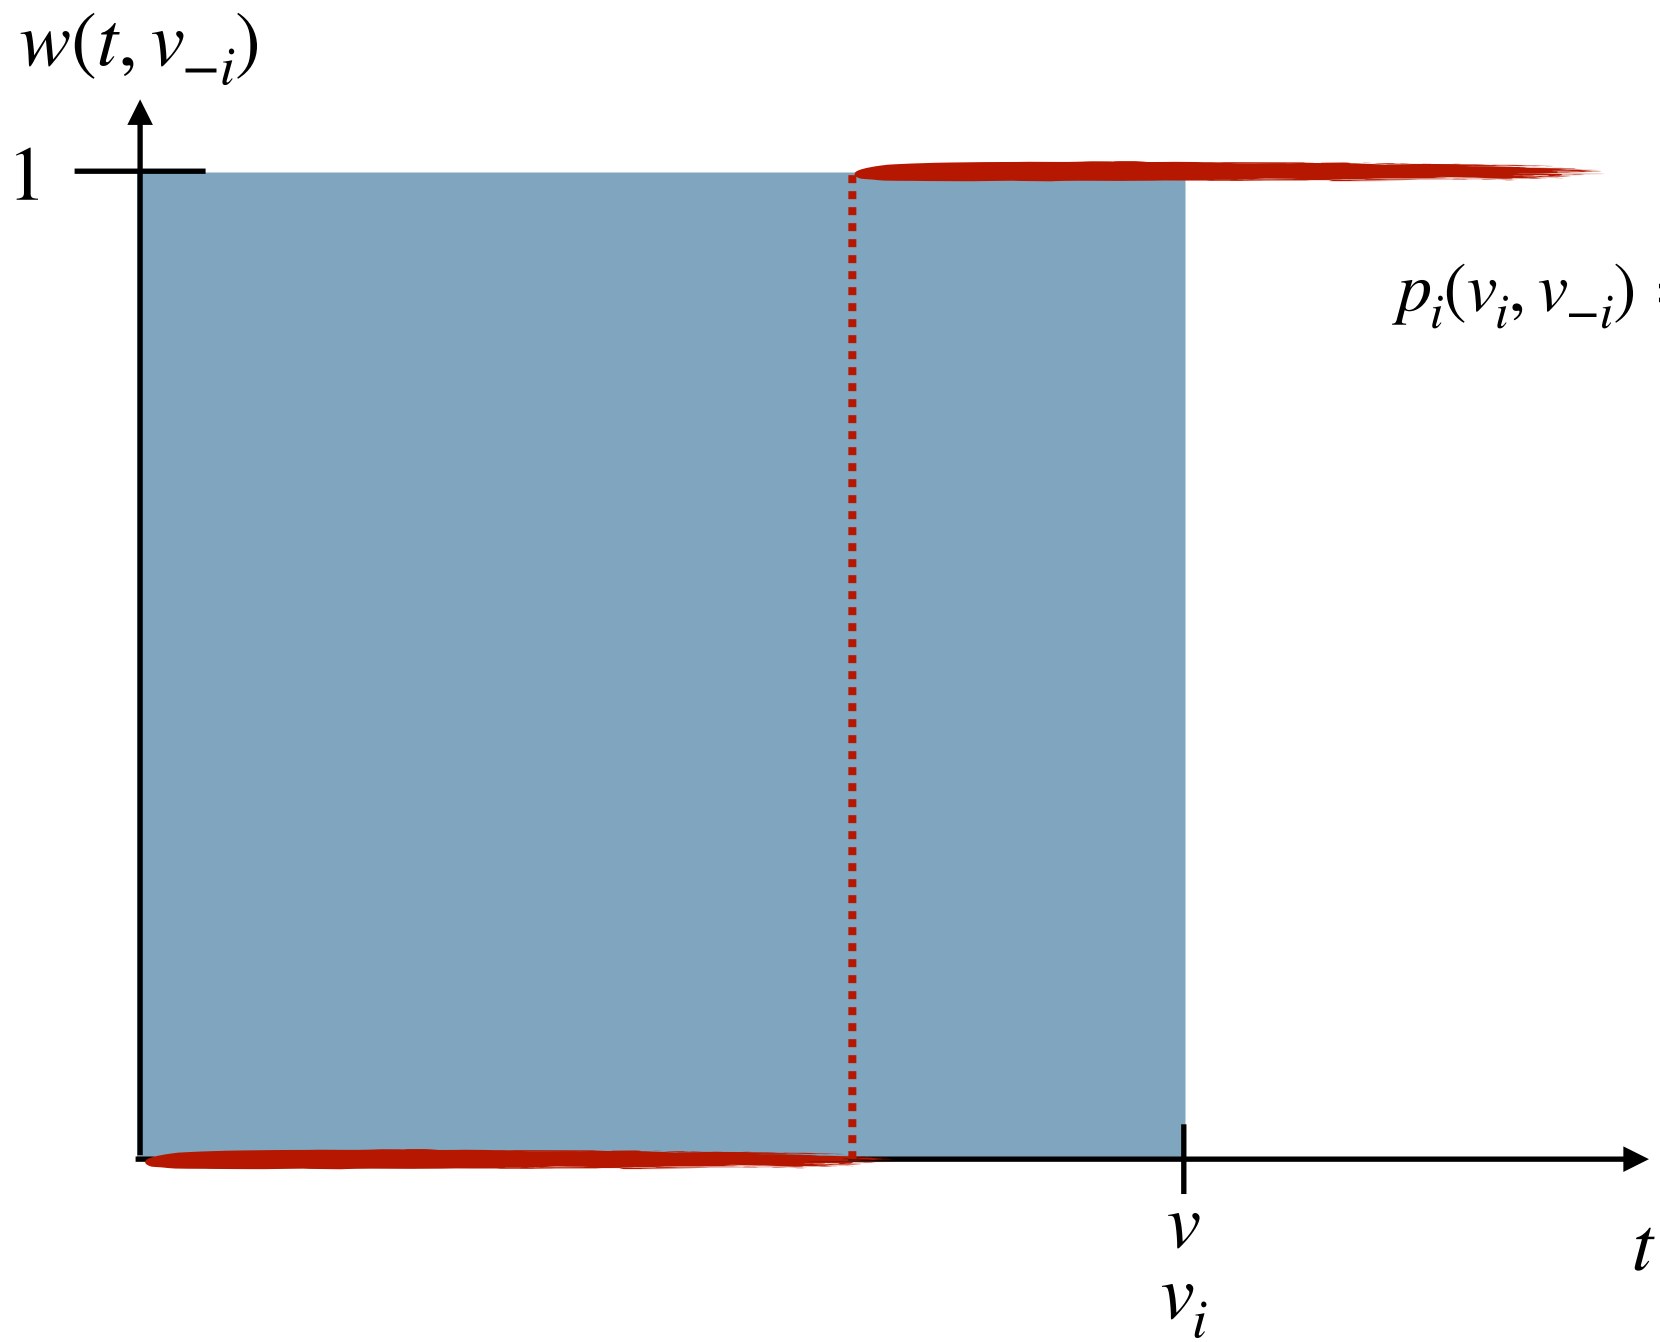

$$p_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})} - \int_{v_i^{\ell}}^{v_i} w(t, v_{-i})dt$$

Some pictures

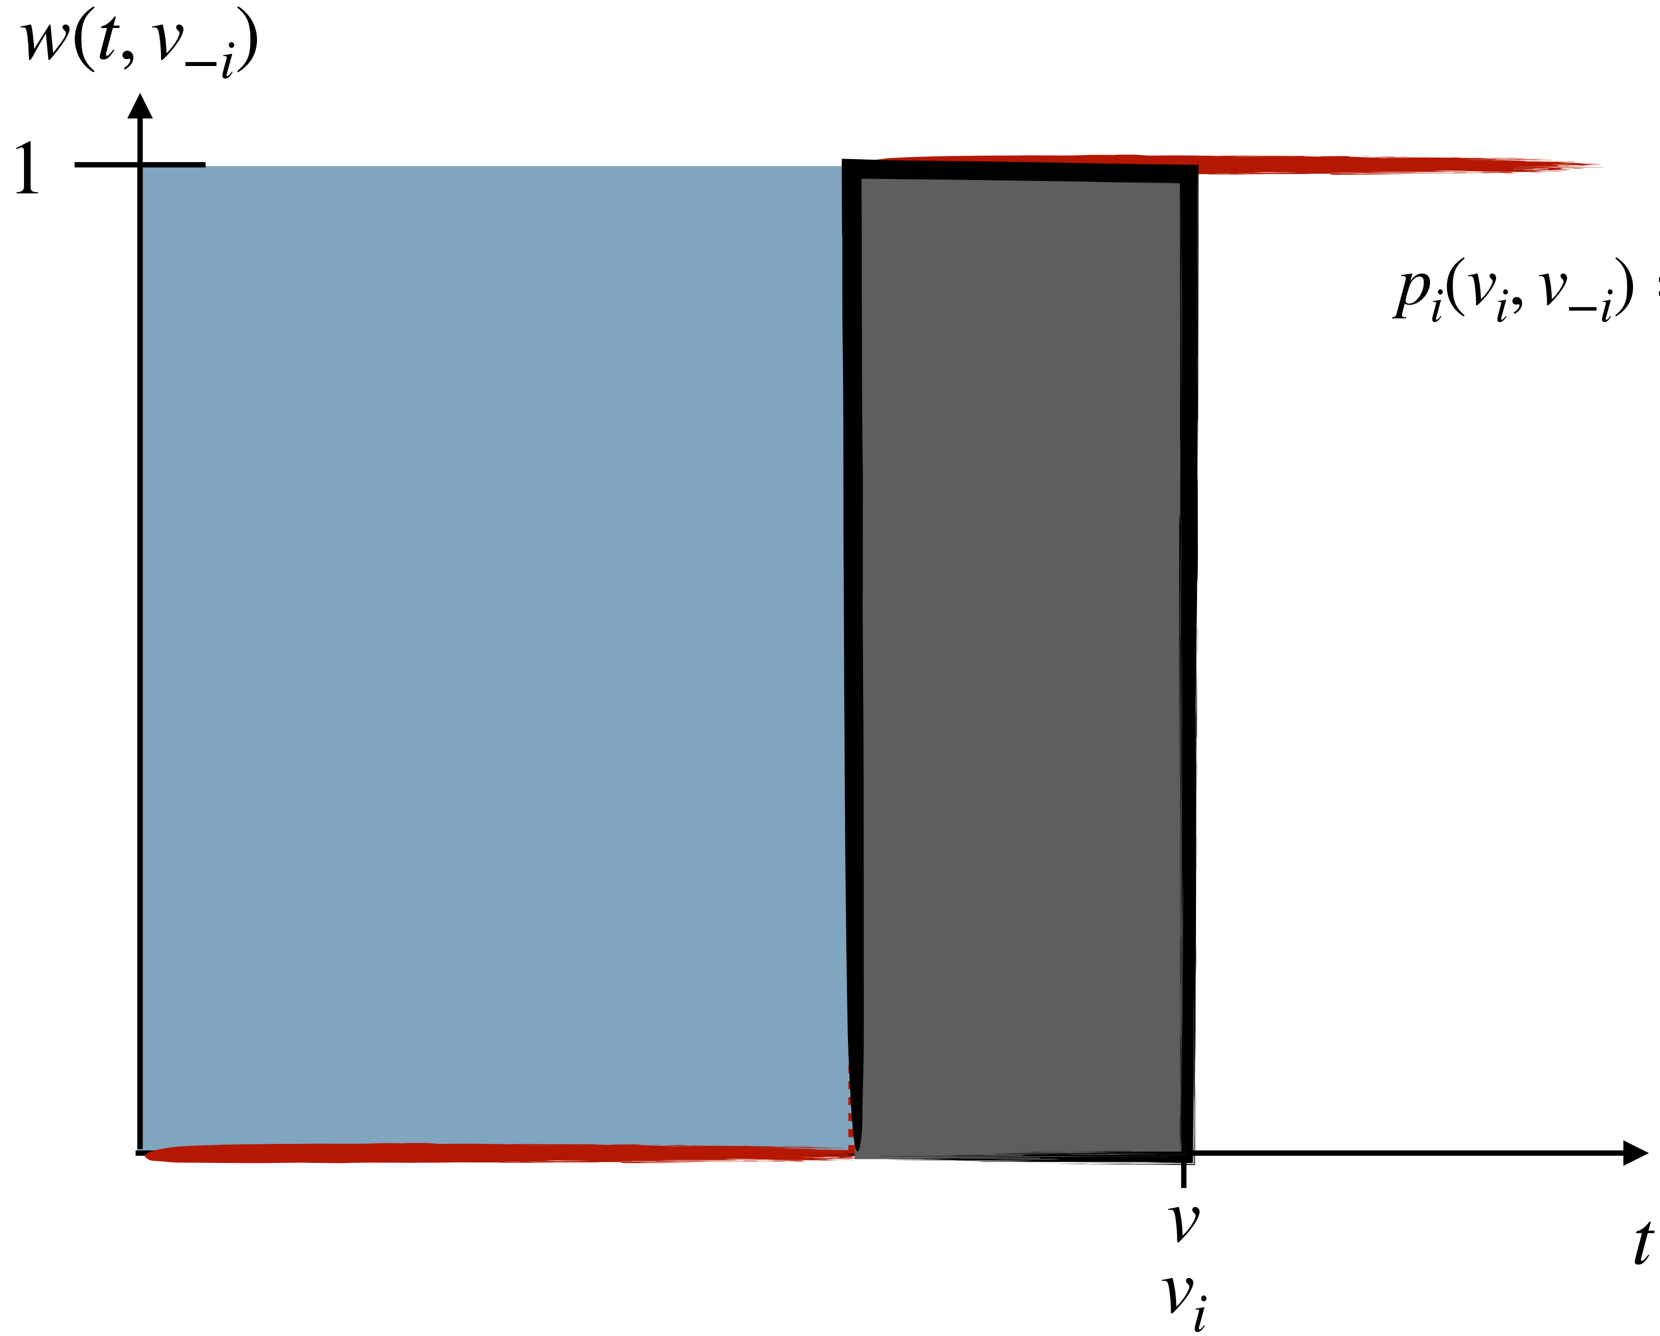

$$p_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})} - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

Some pictures

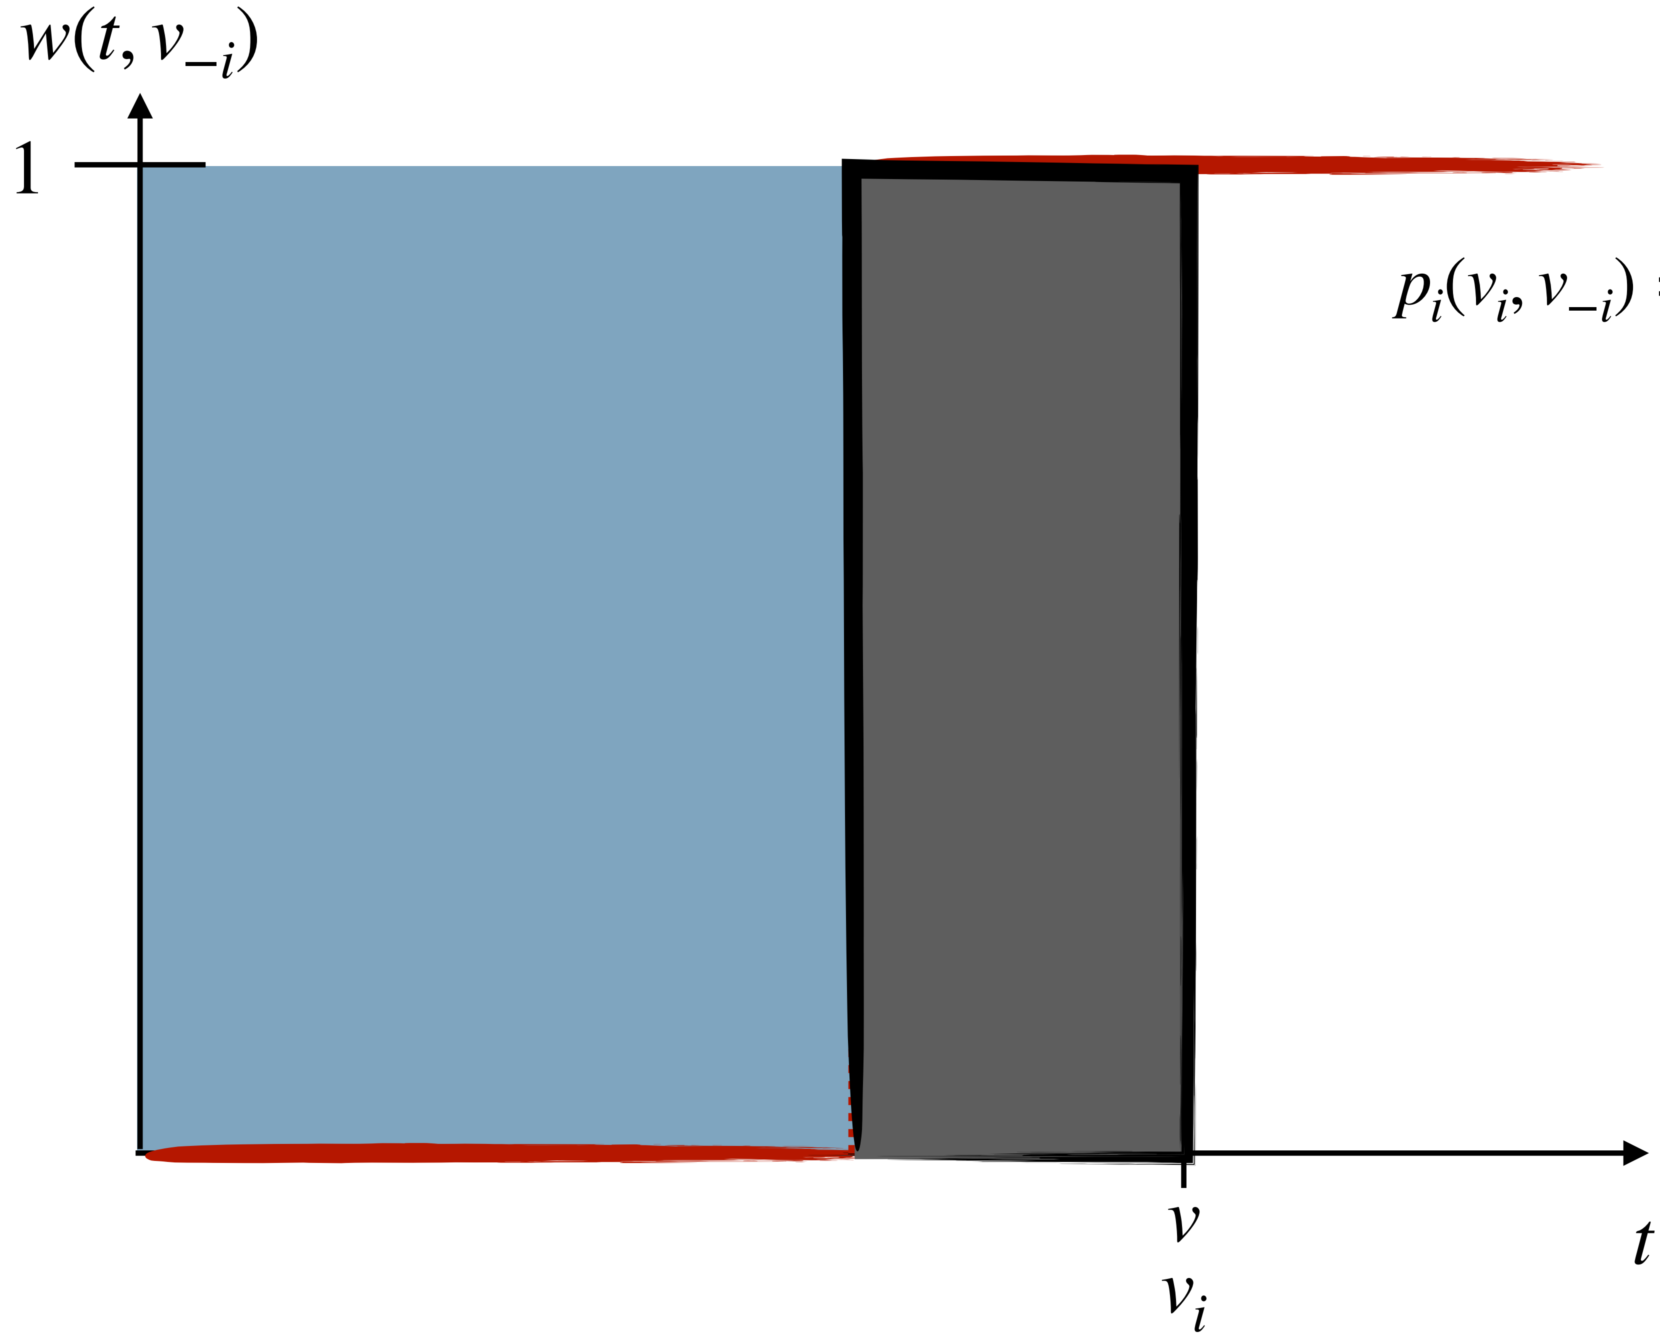

$$p_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})} - \underbrace{\int_{v_i^\ell}^{v_i} w(t, v_{-i})dt}$$

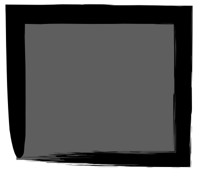

Some pictures

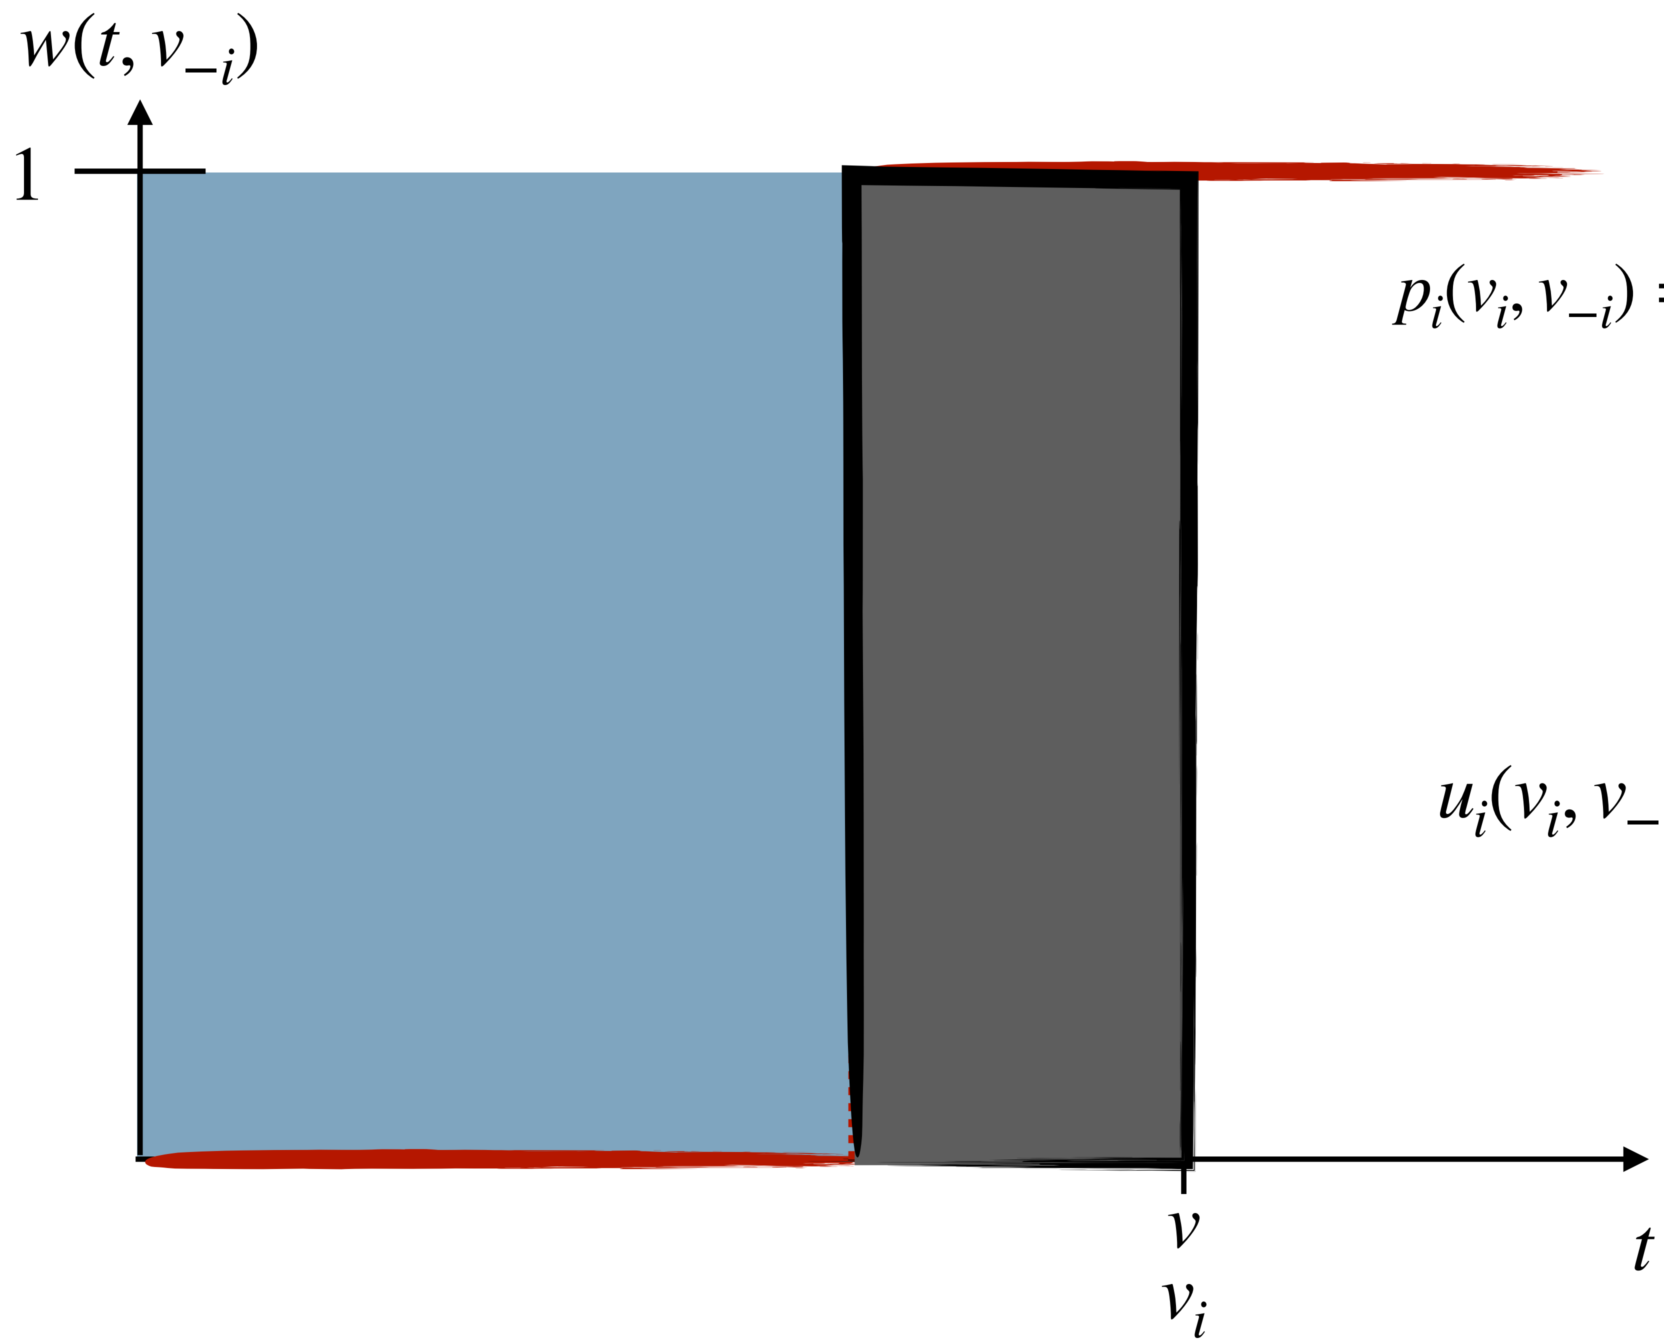

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i}) dt$$

$$u_i(v_i, v_{-i}) = v \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

Some pictures

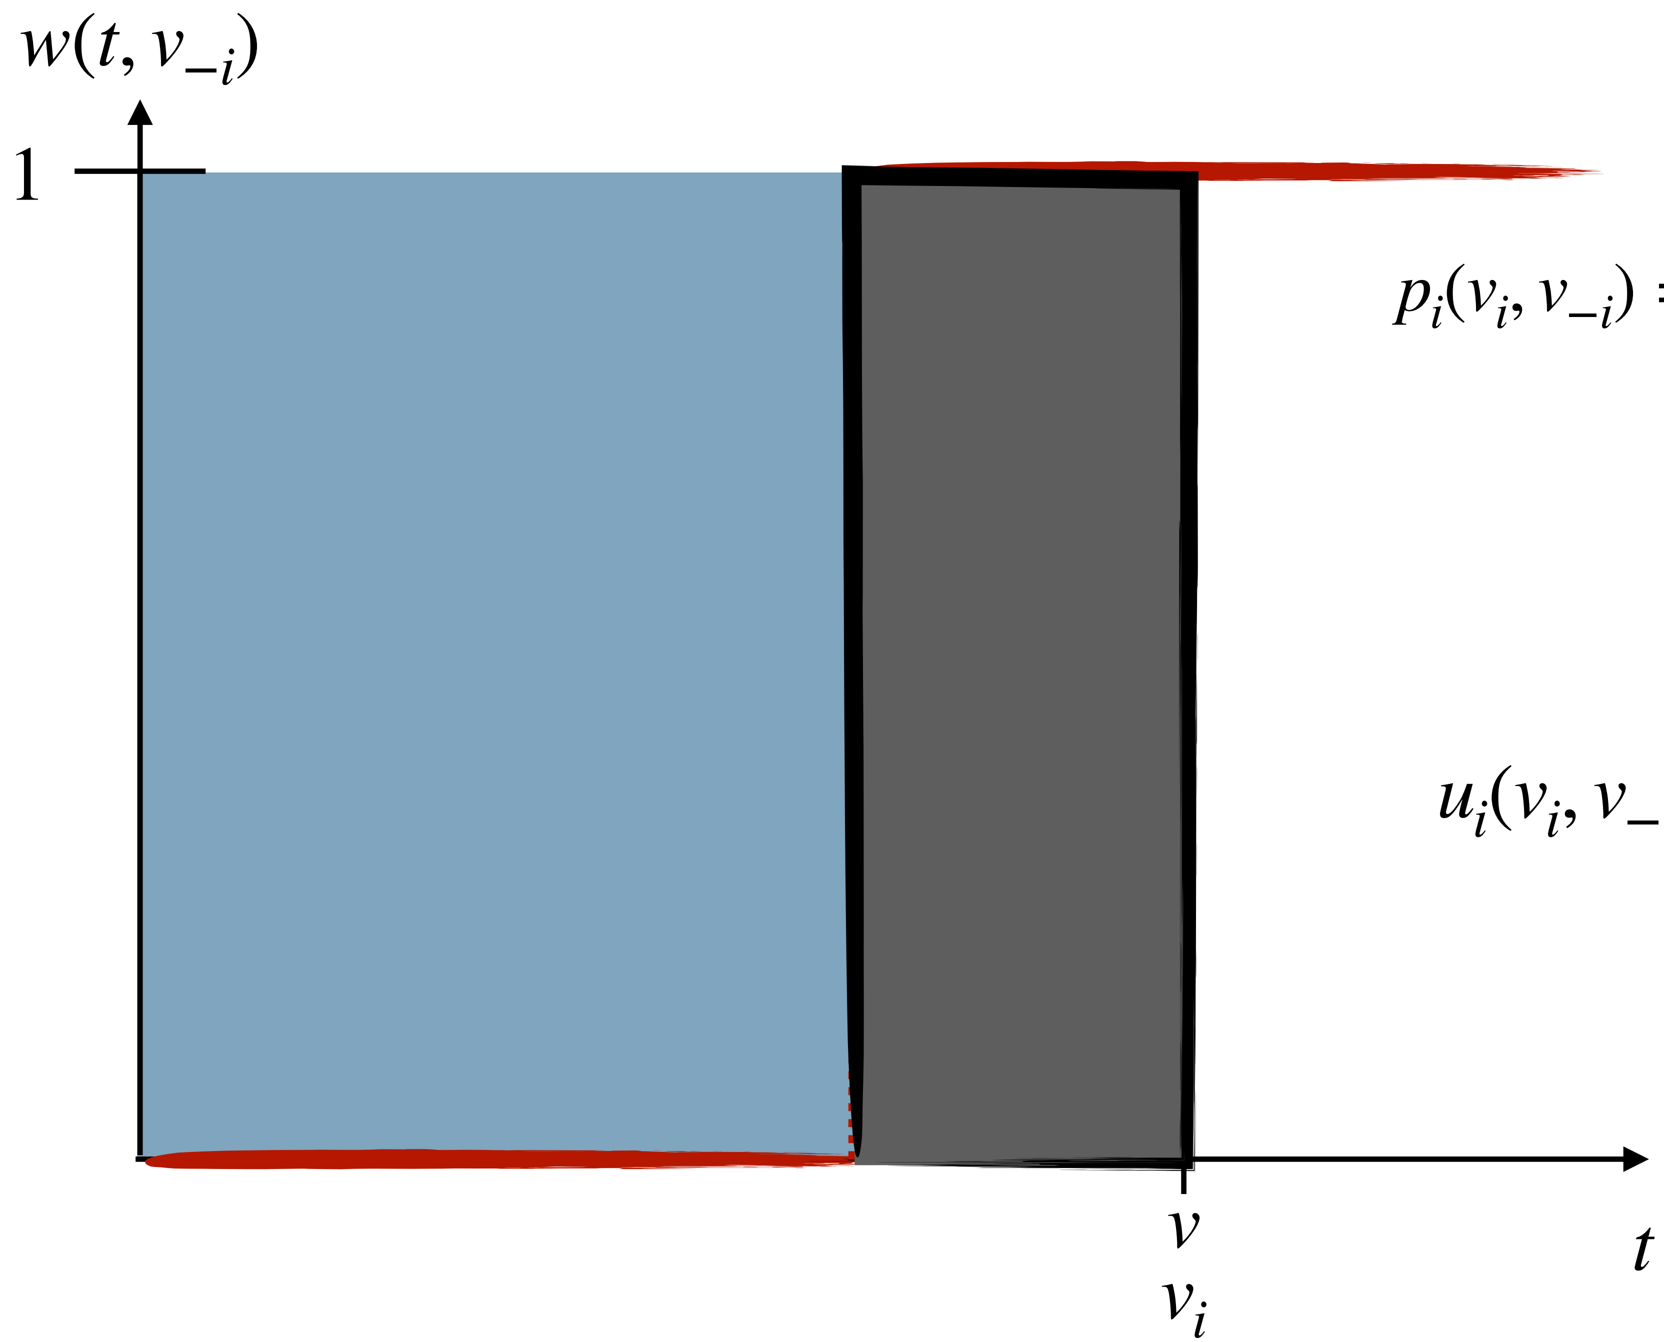

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^{\ell}}^{v_i} w(t, v_{-i})dt$$

$$u_i(v_i, v_{-i}) = v \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

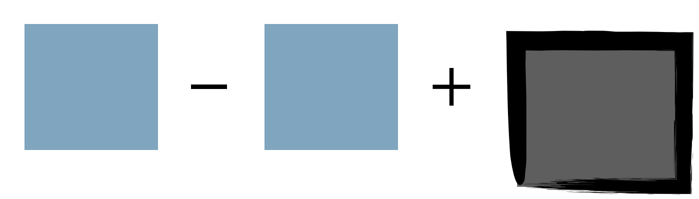

Some pictures

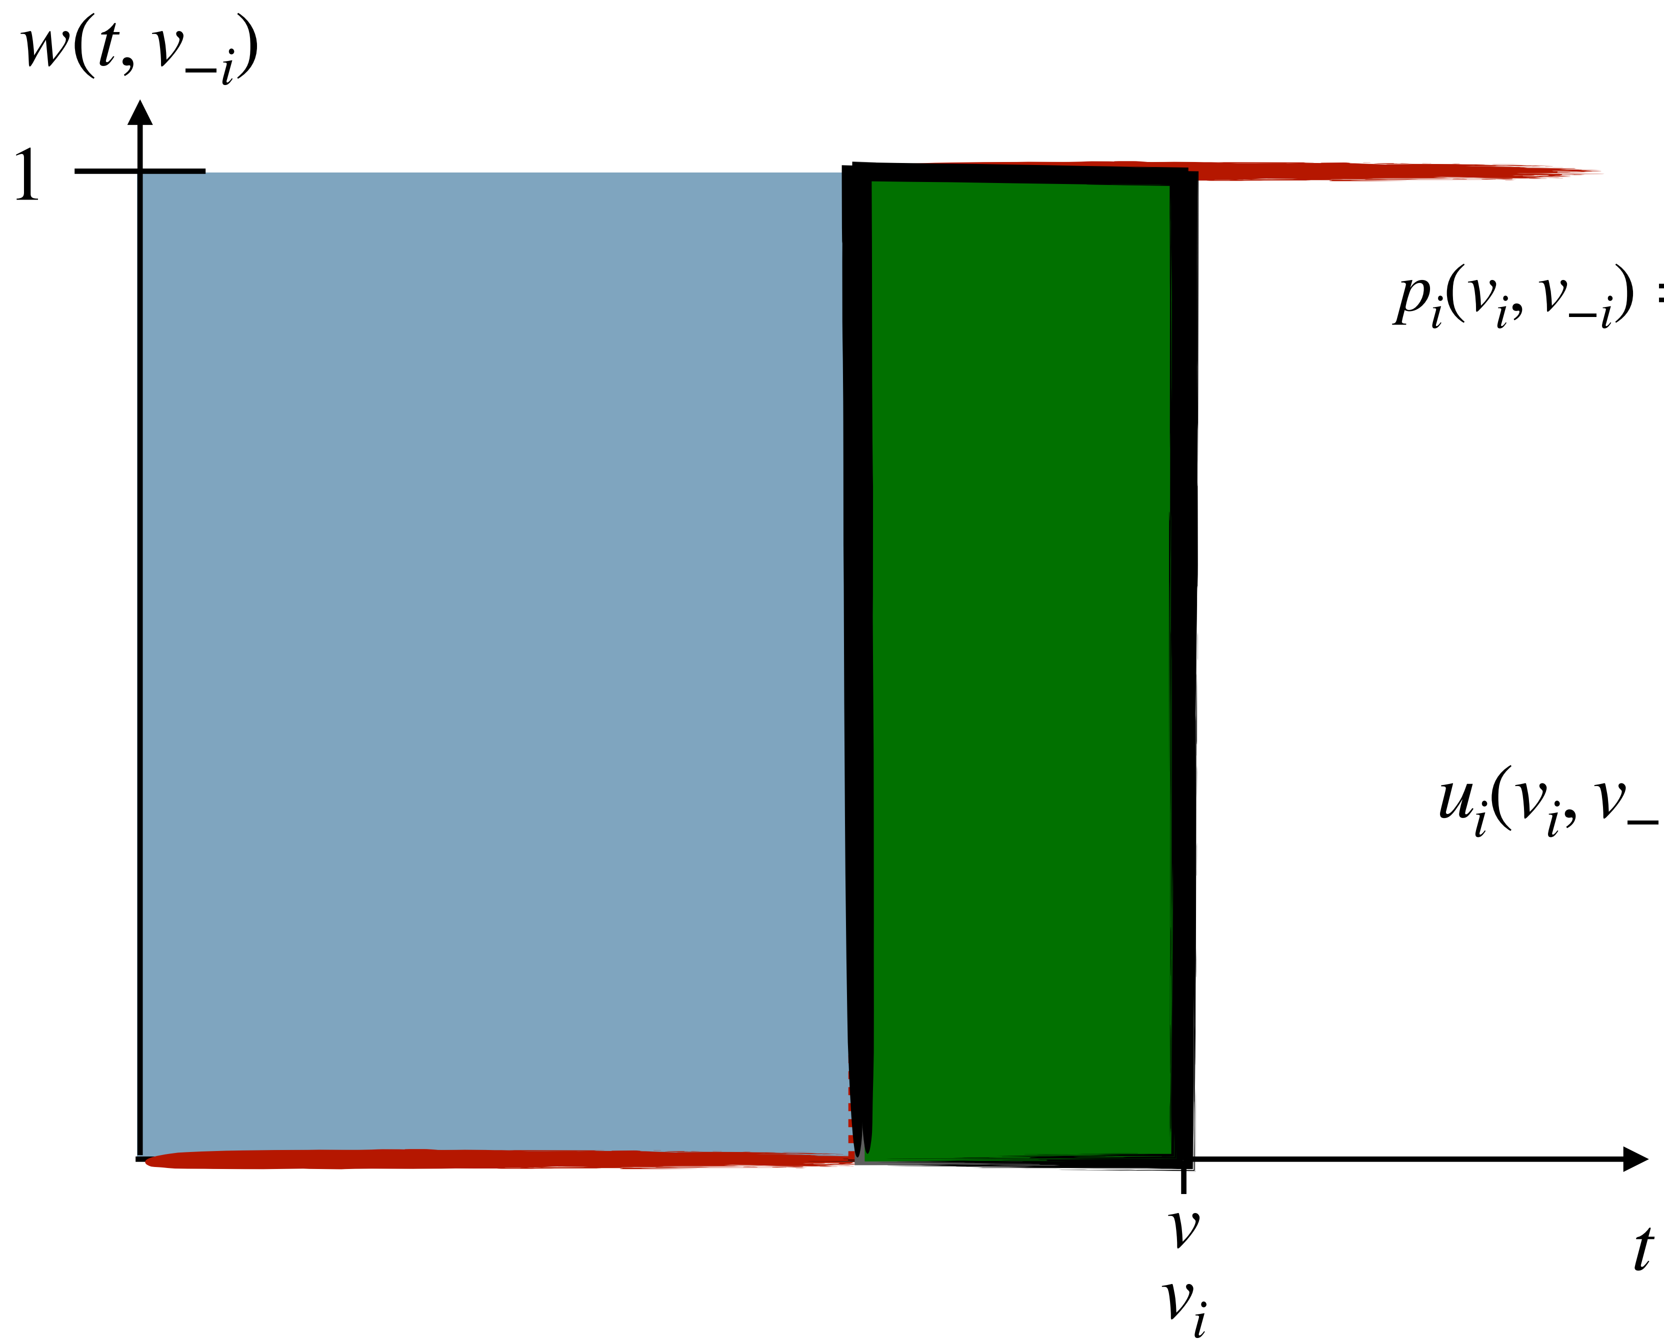

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^{\ell}}^{v_i} w(t, v_{-i})dt$$

$$u_i(v_i, v_{-i}) = v \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

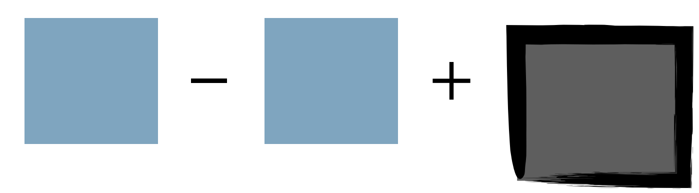

Some pictures

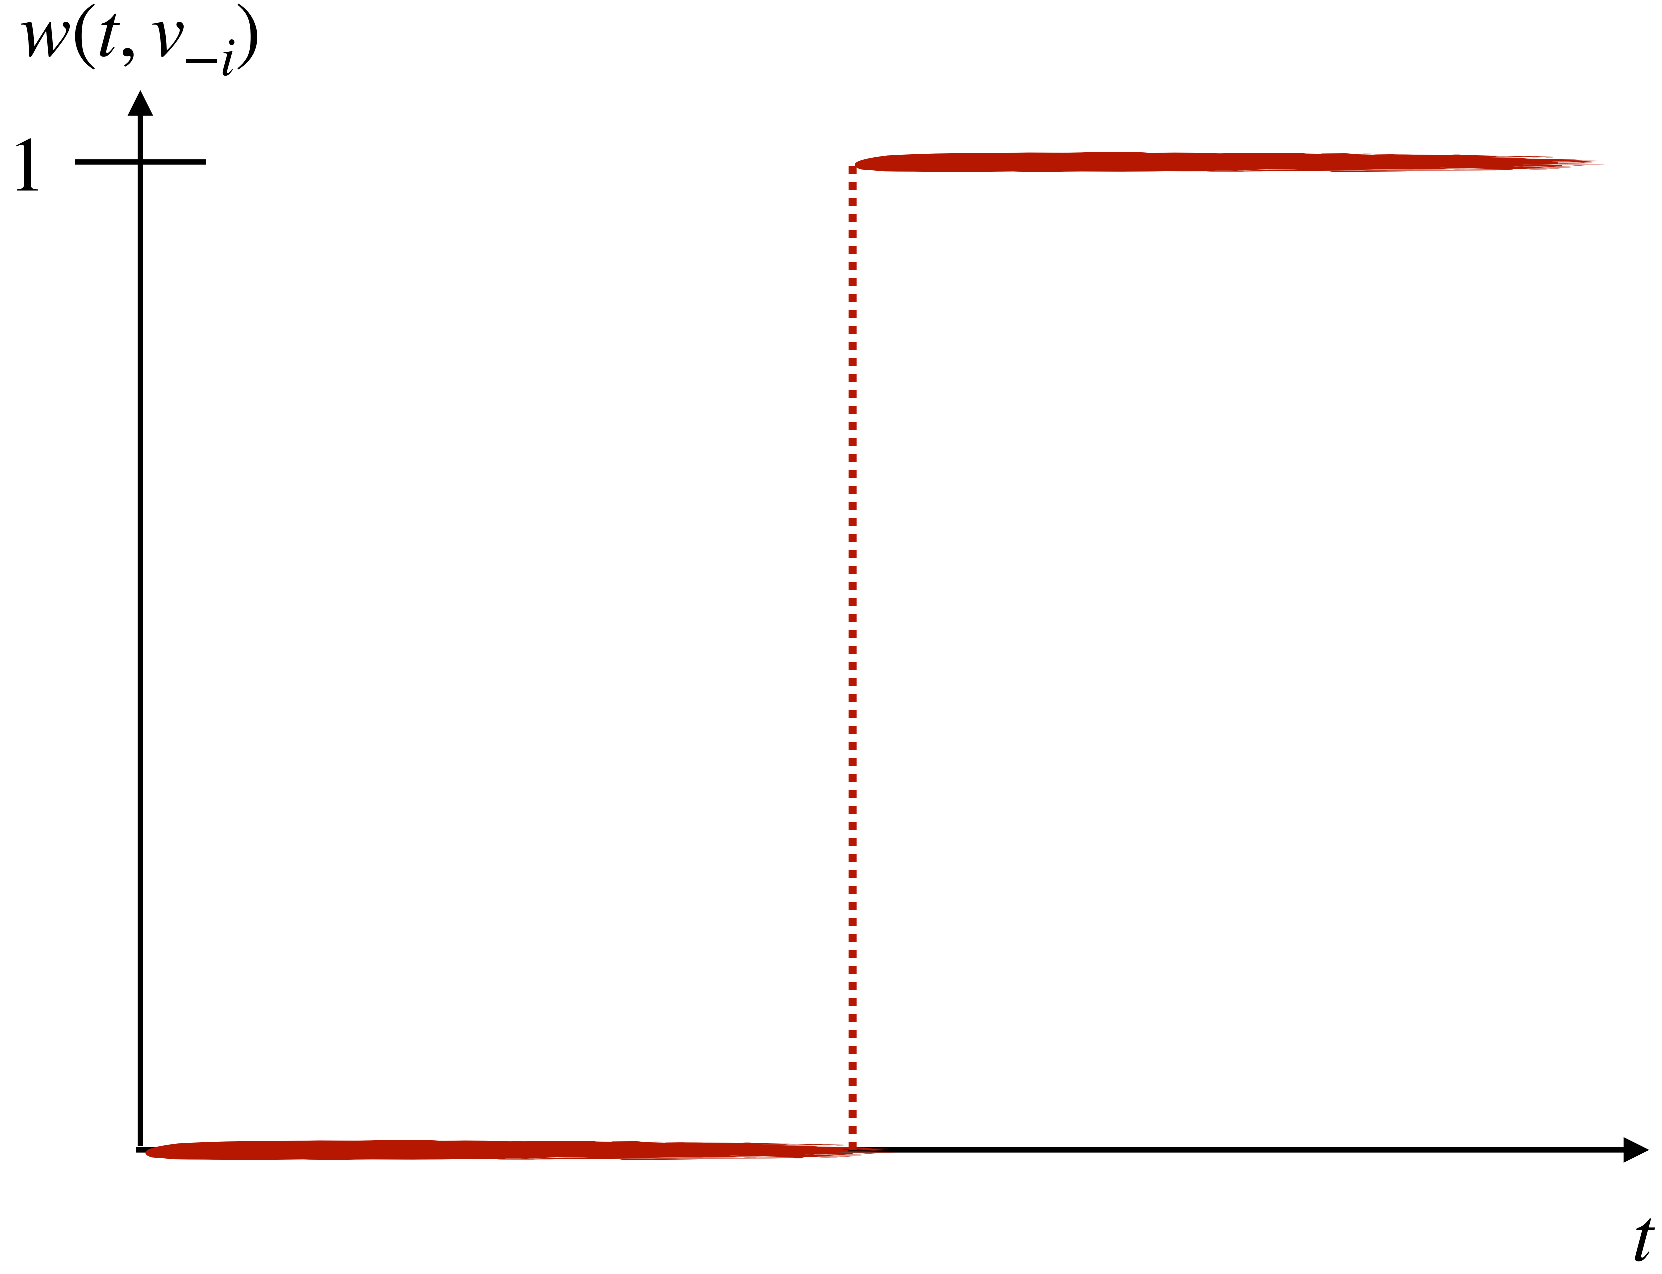

Some pictures

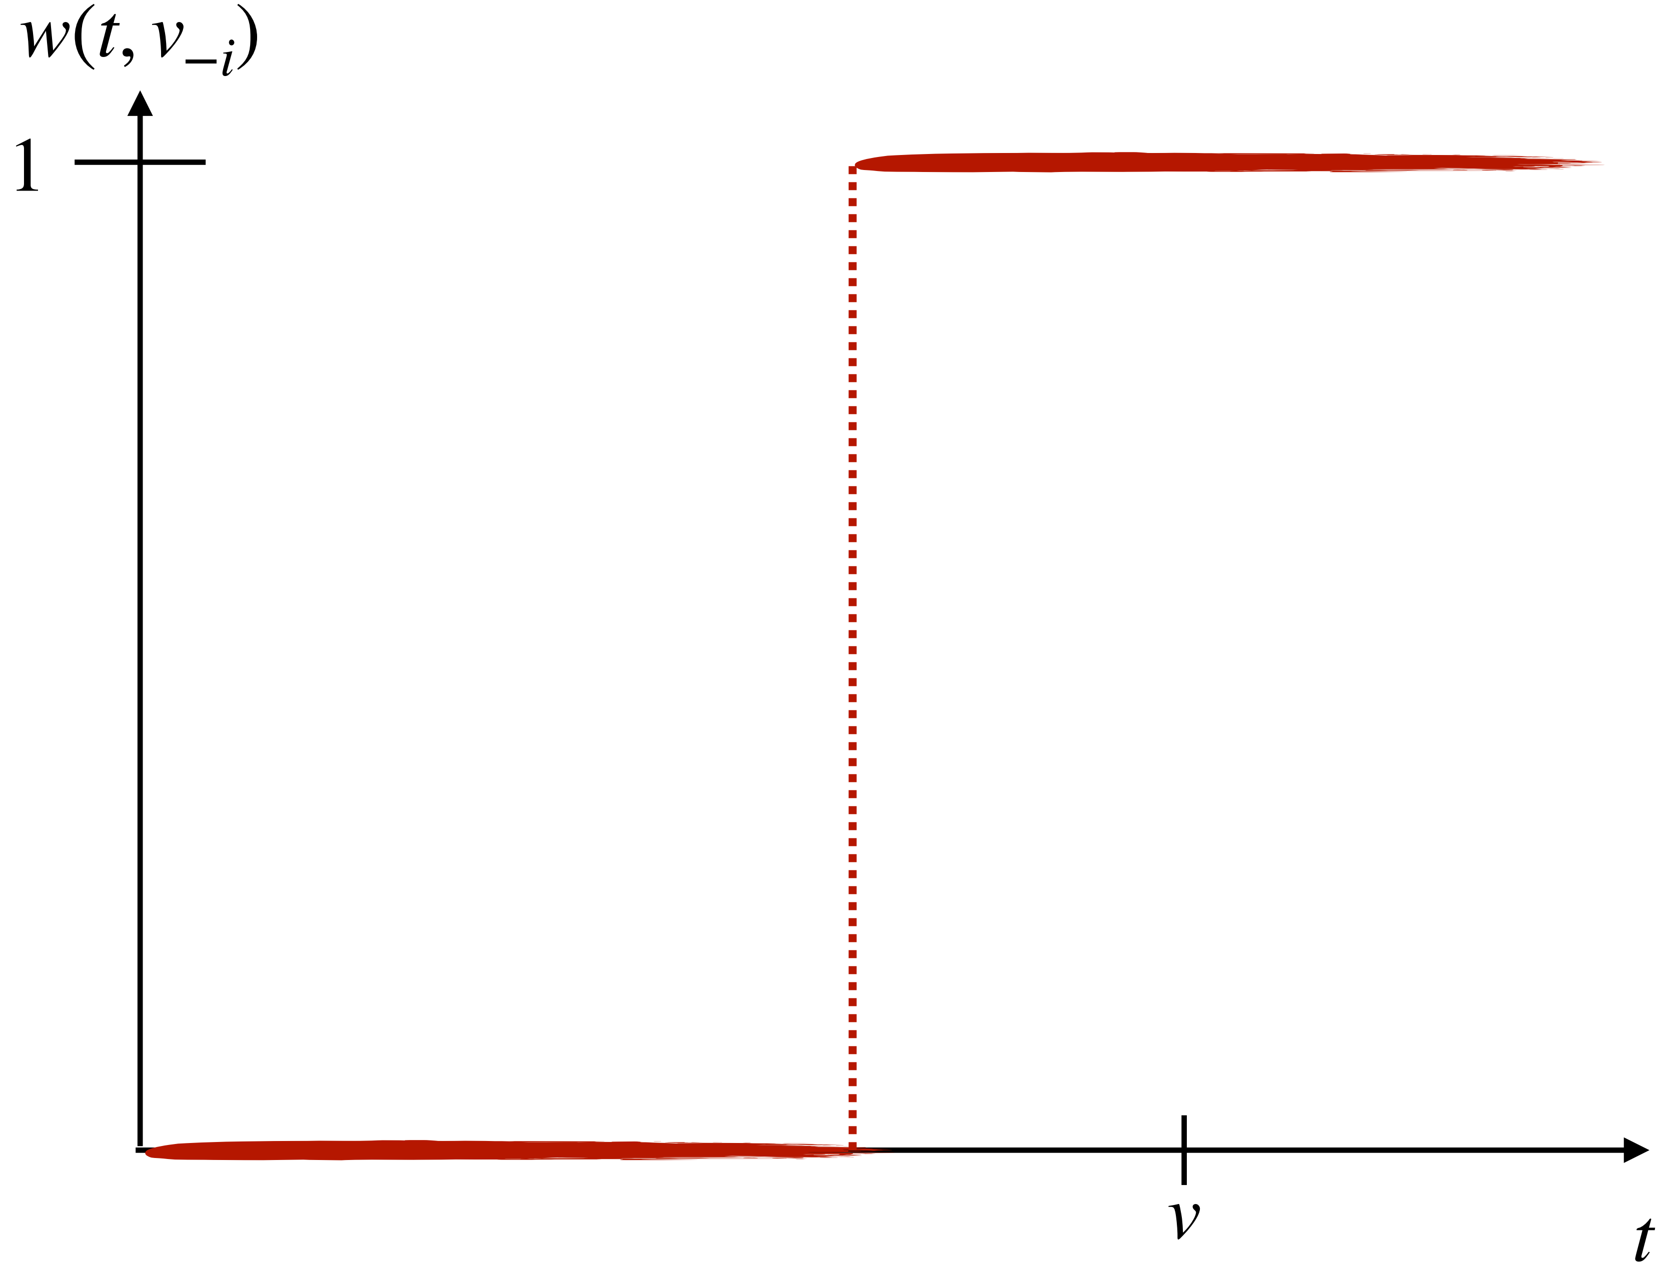

Some pictures

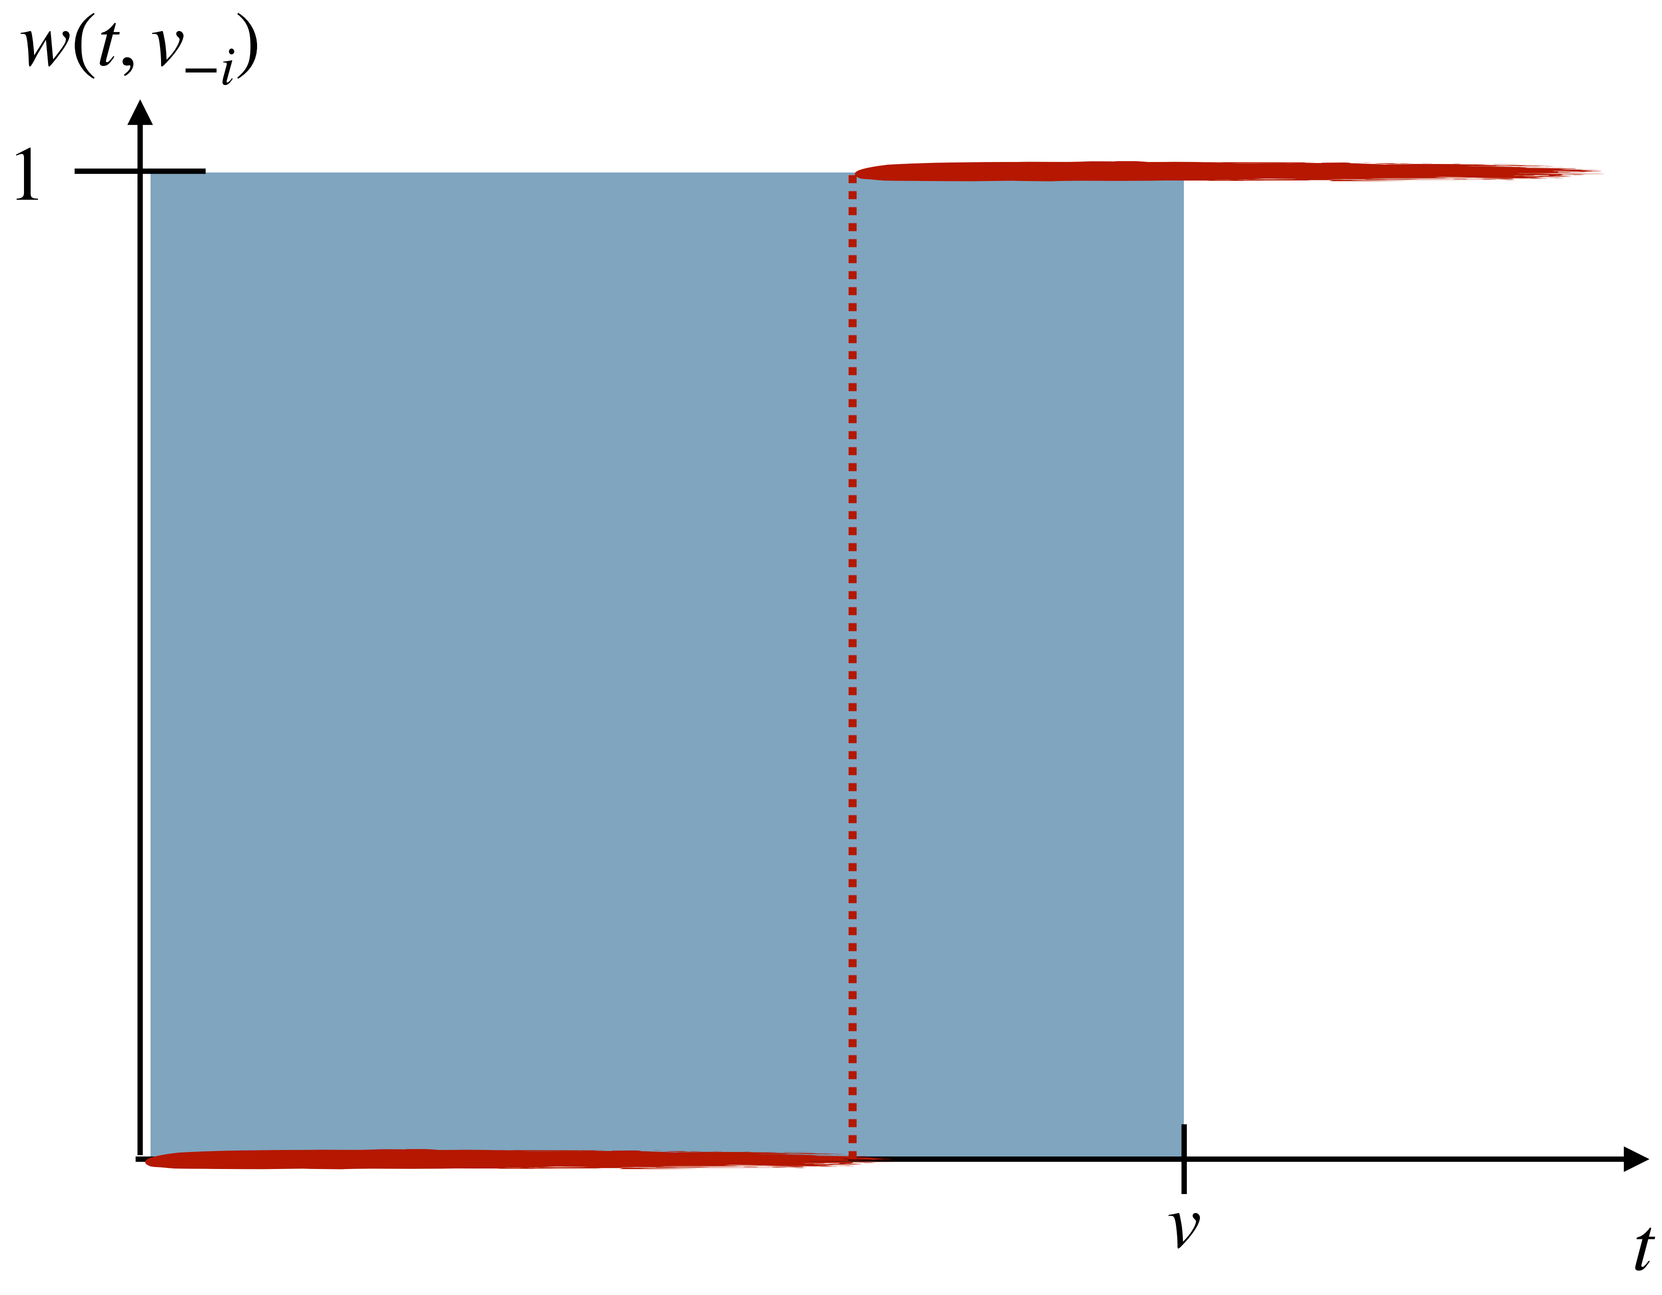

Some pictures

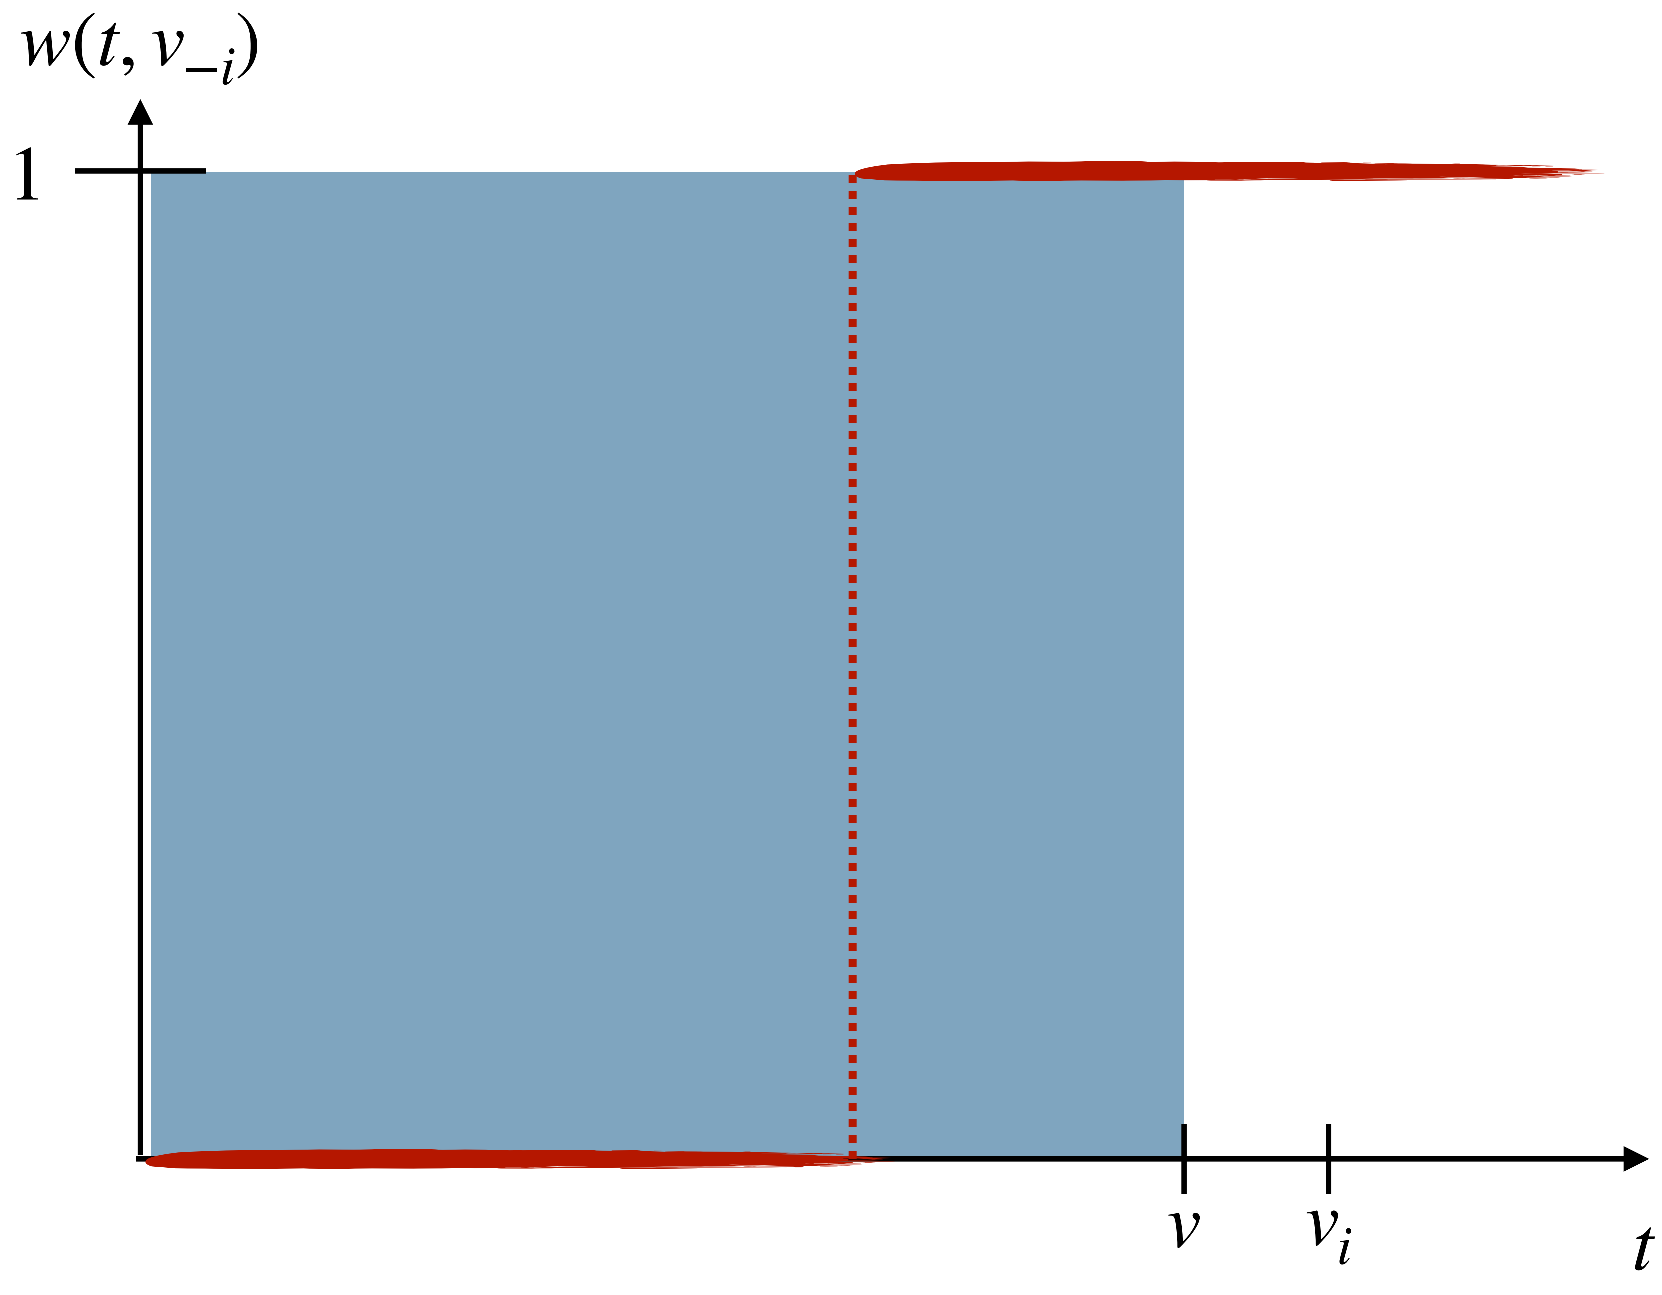

Some pictures

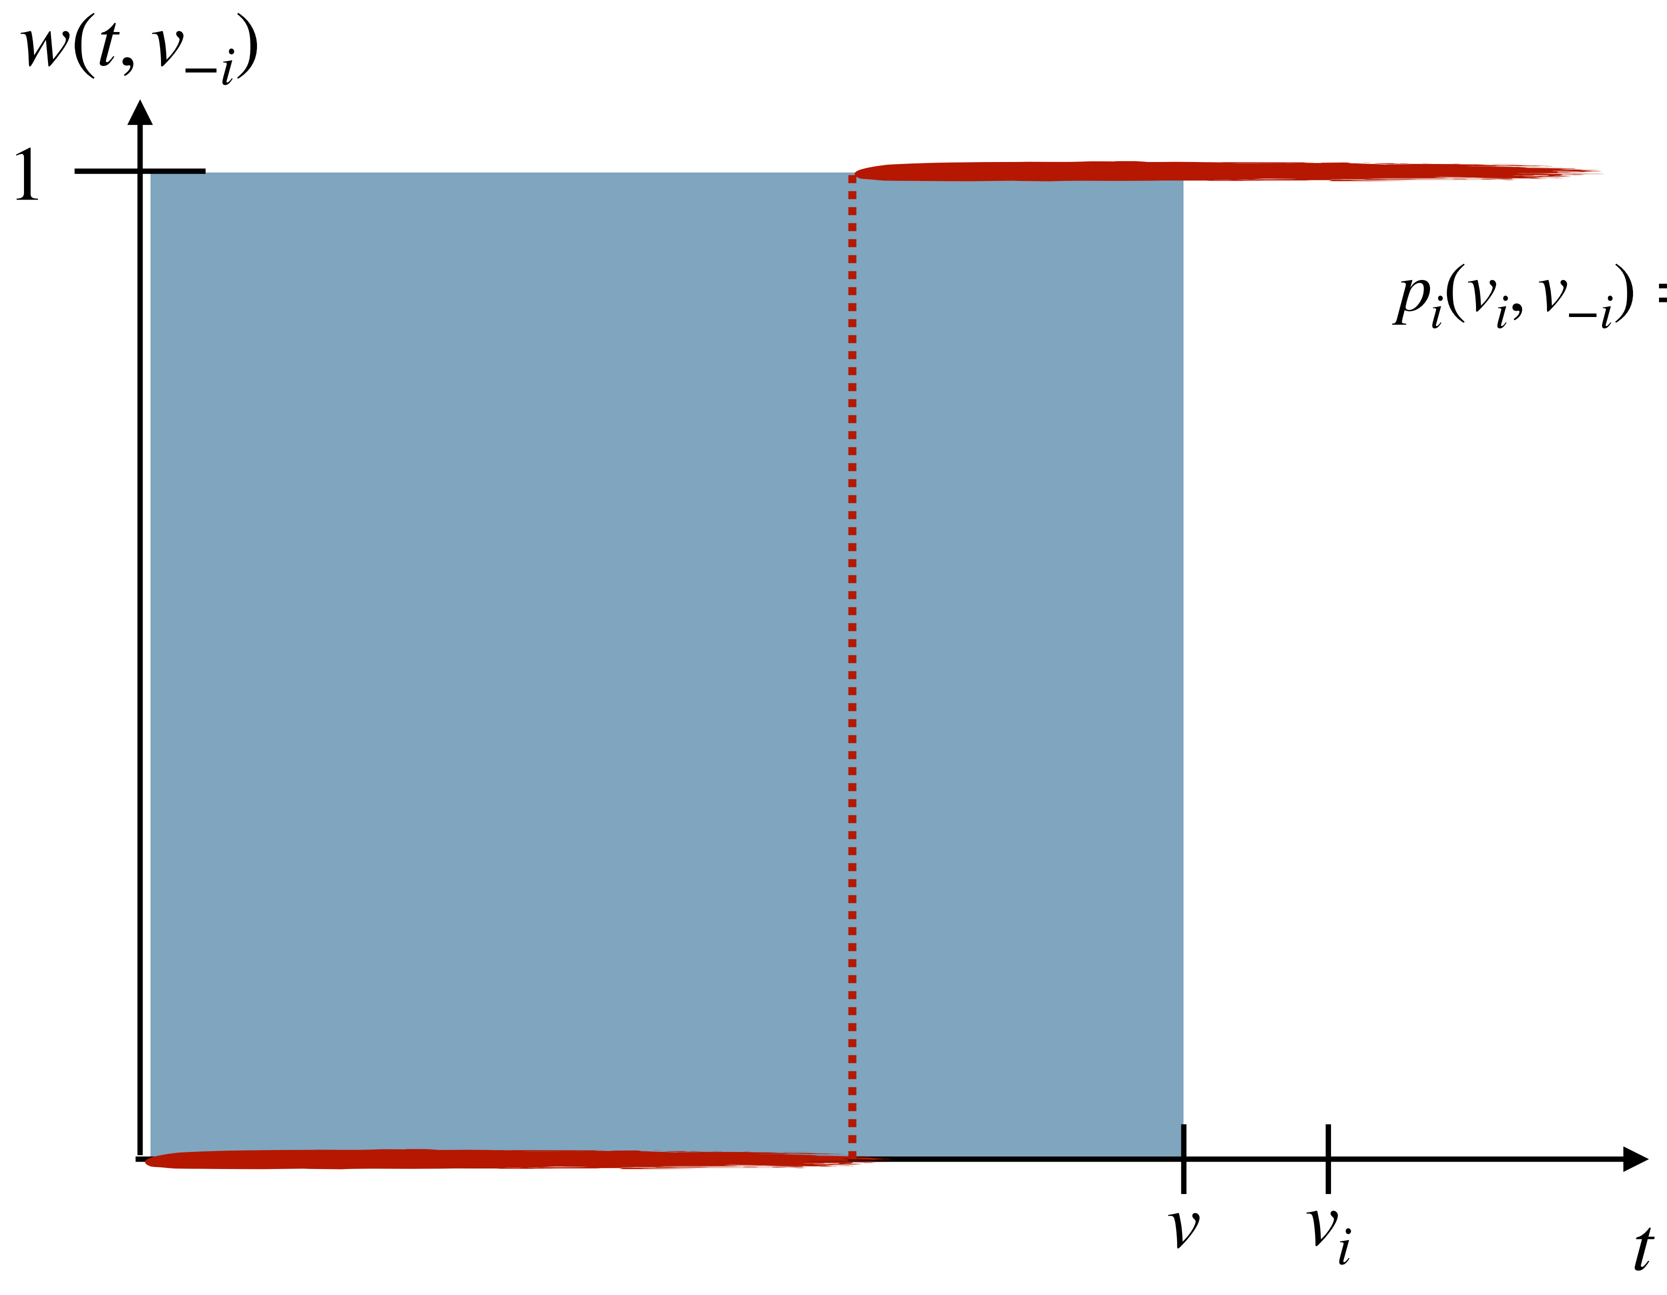

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^{\ell}}^{v_i} w(t, v_{-i})dt$$

Some pictures

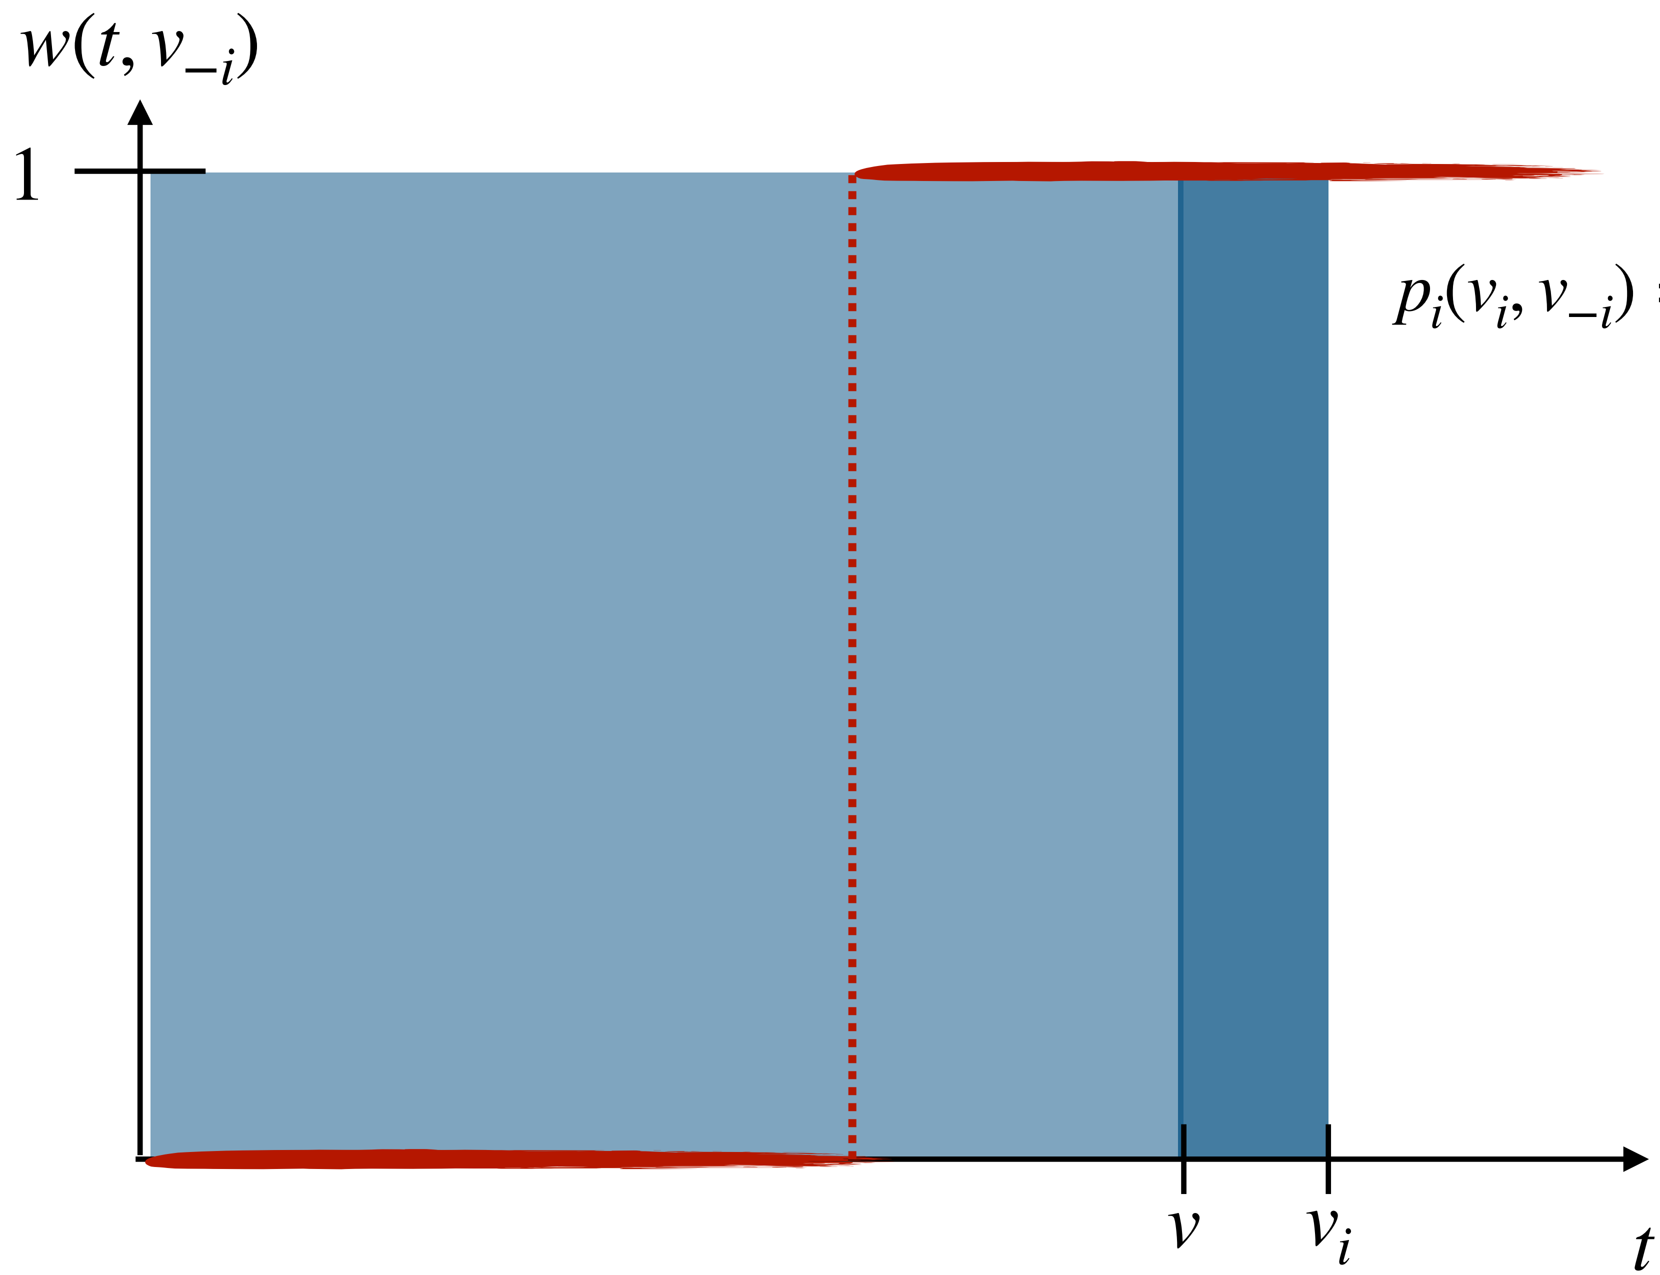

$$p_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})} - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

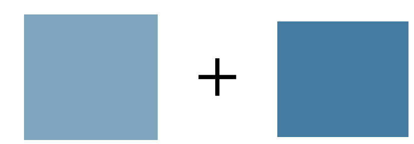

Some pictures

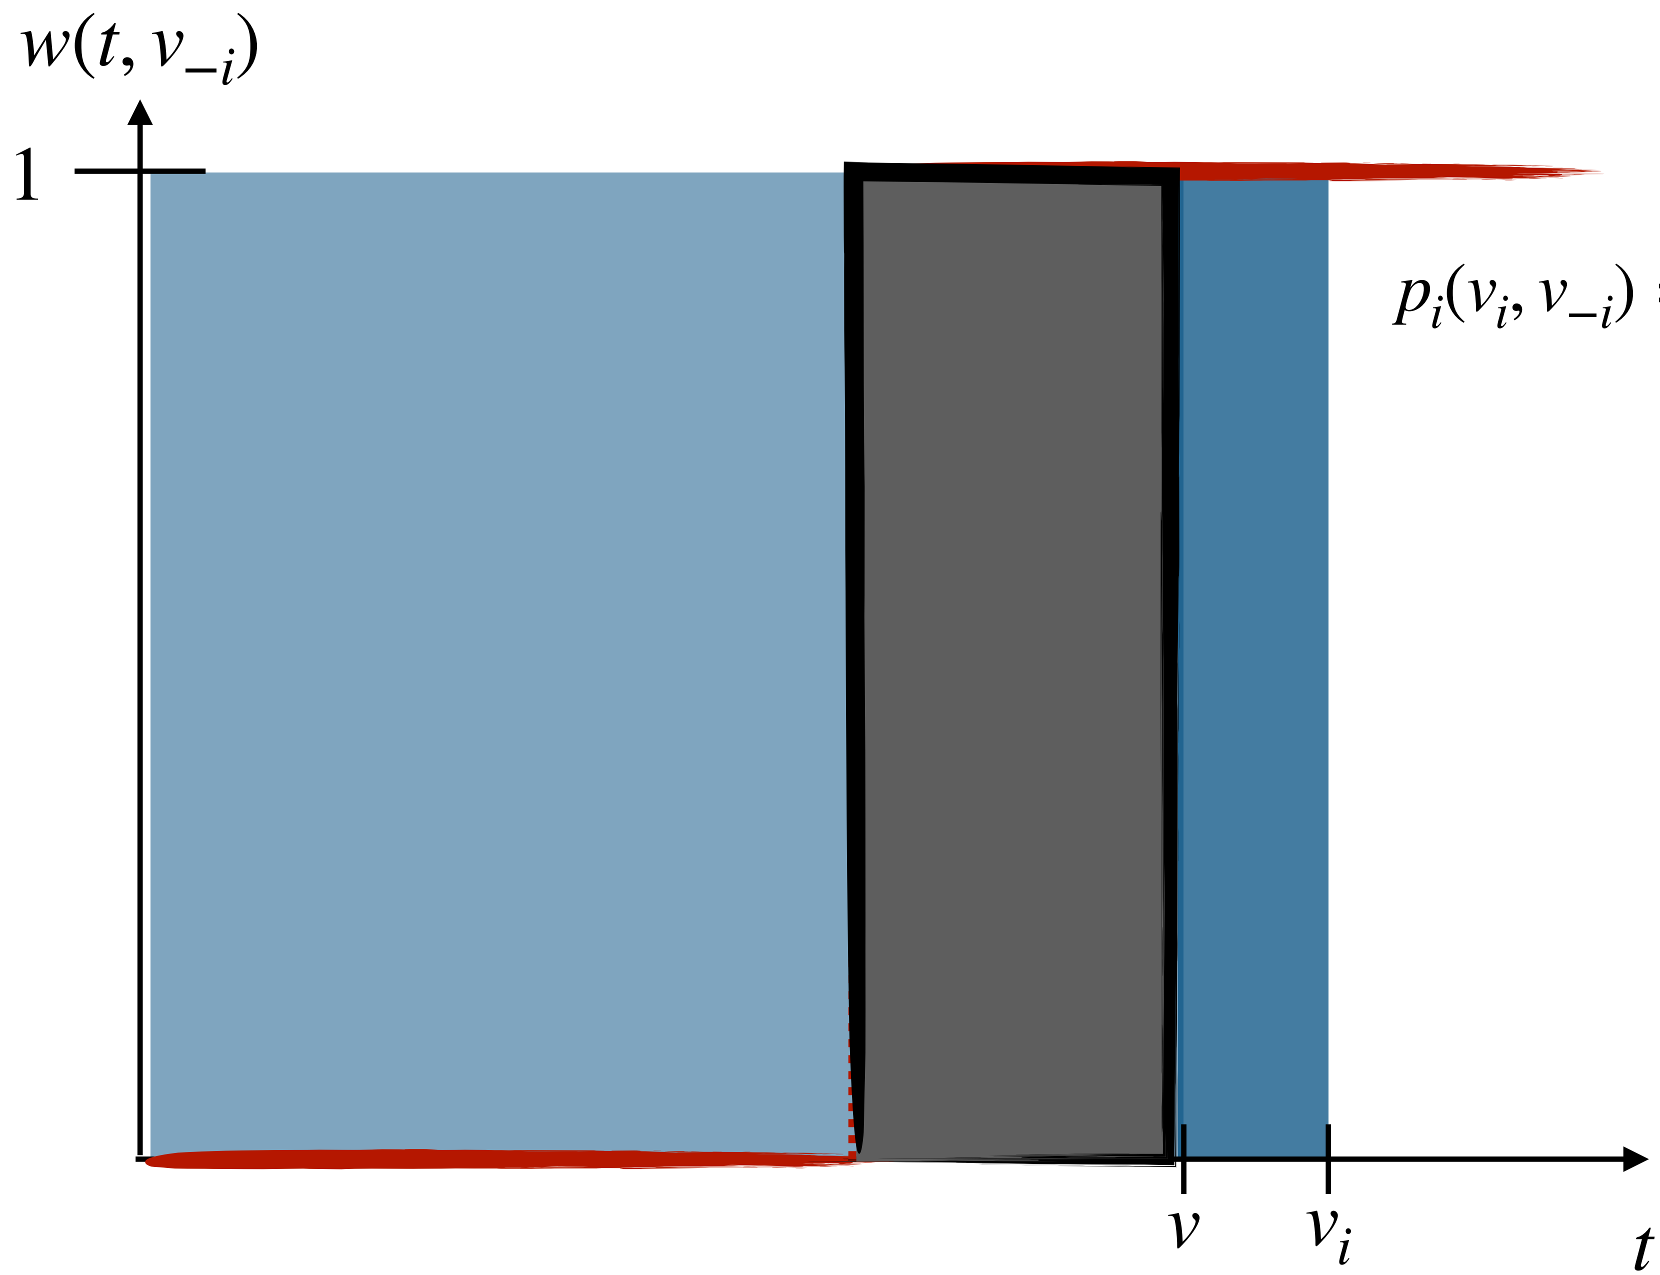

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

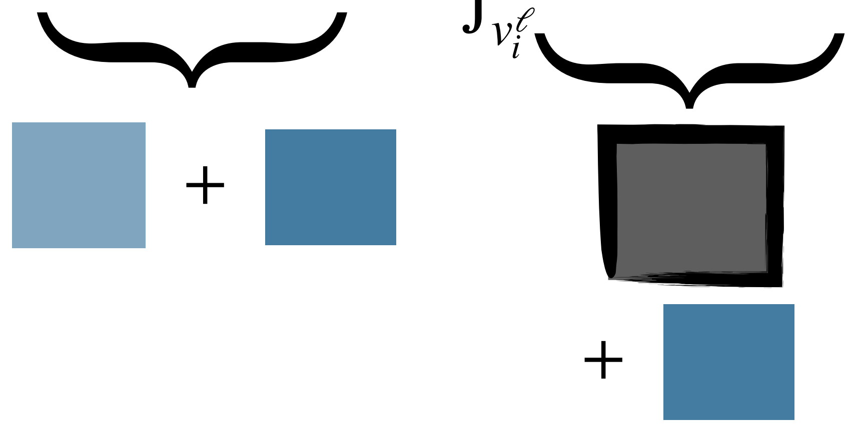

Some pictures

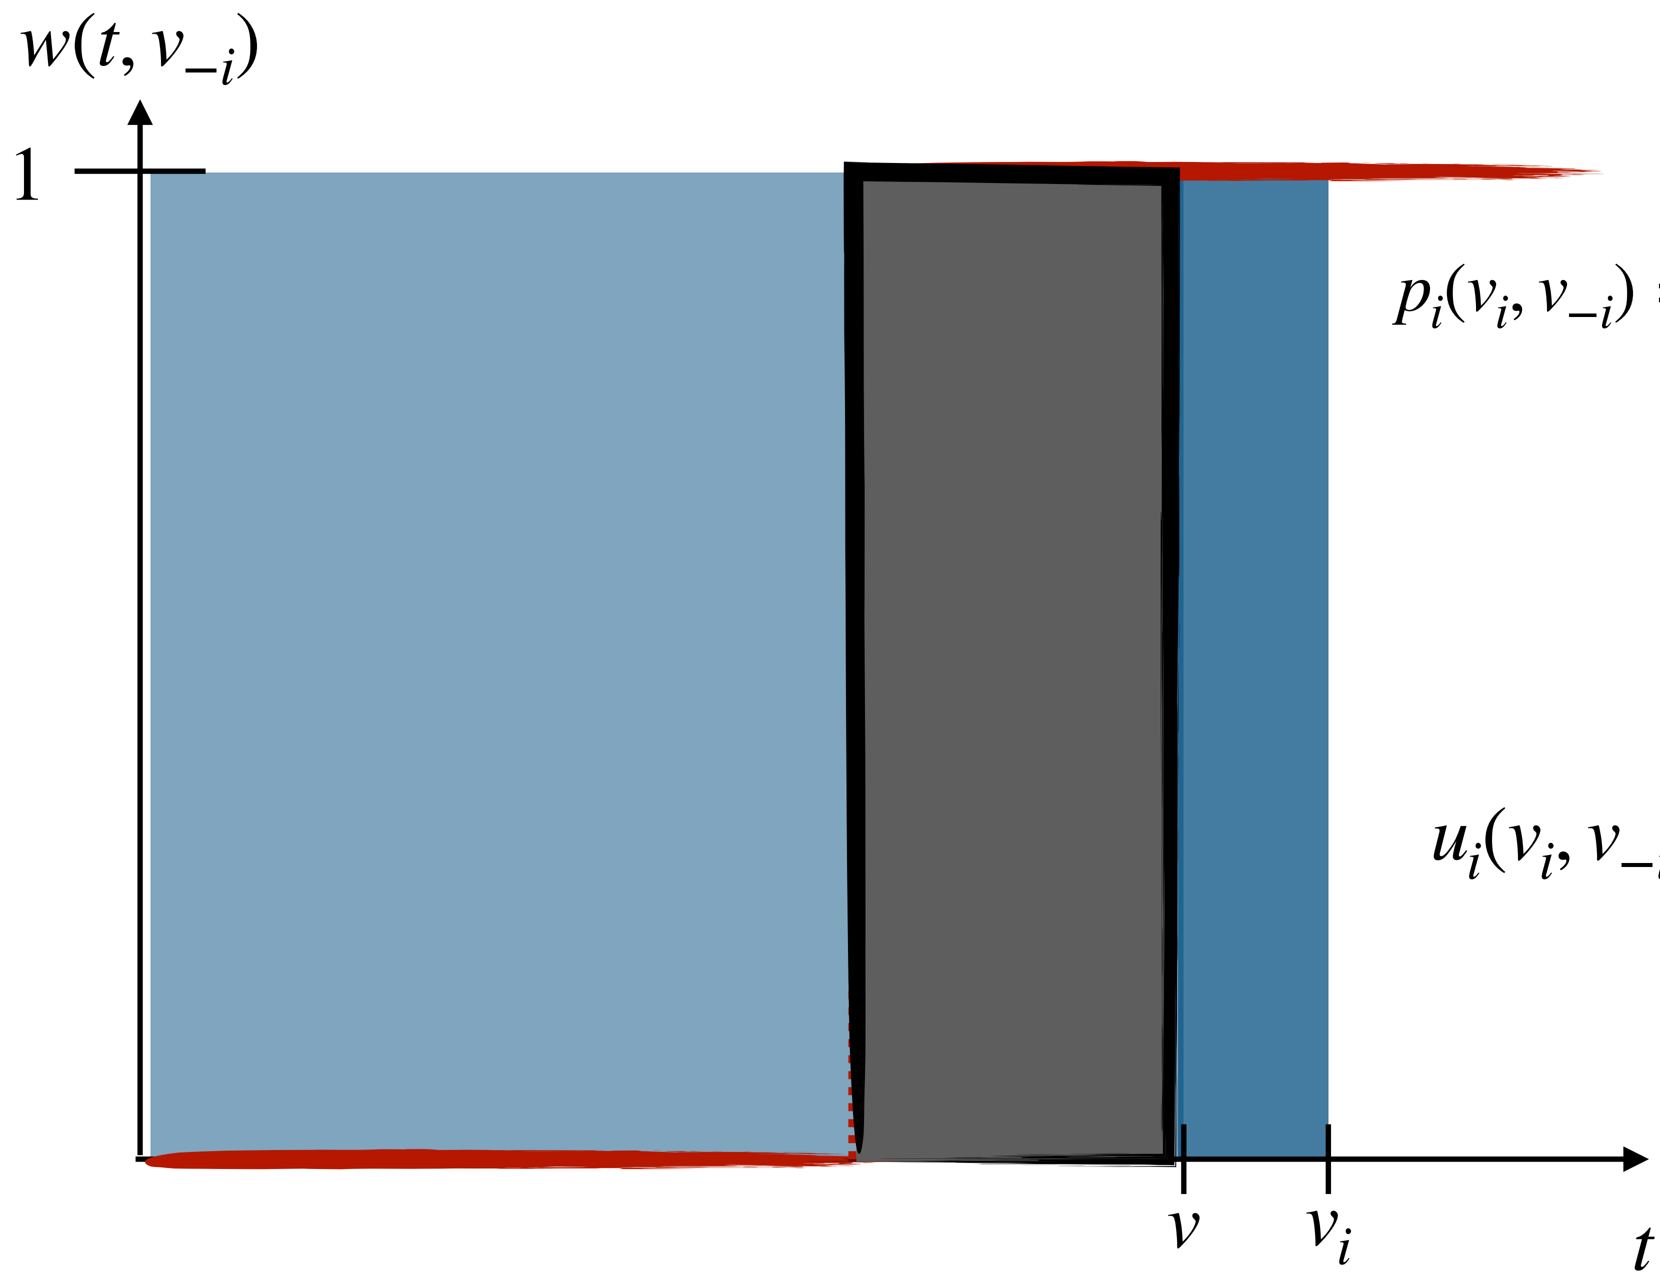

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

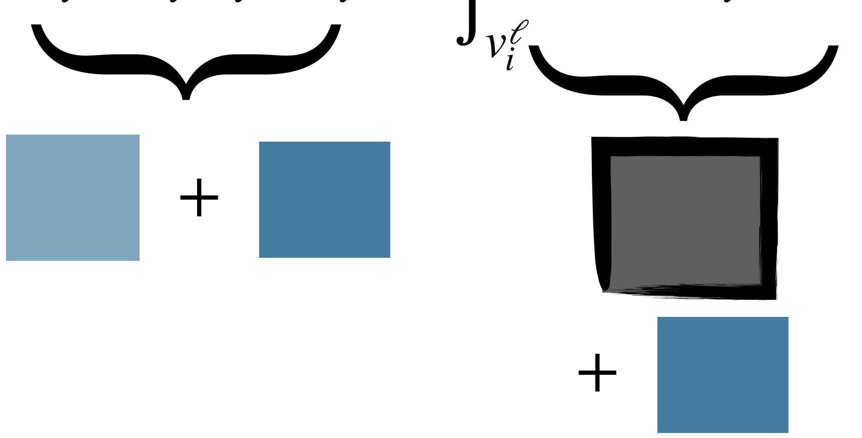

$$u_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

Some pictures

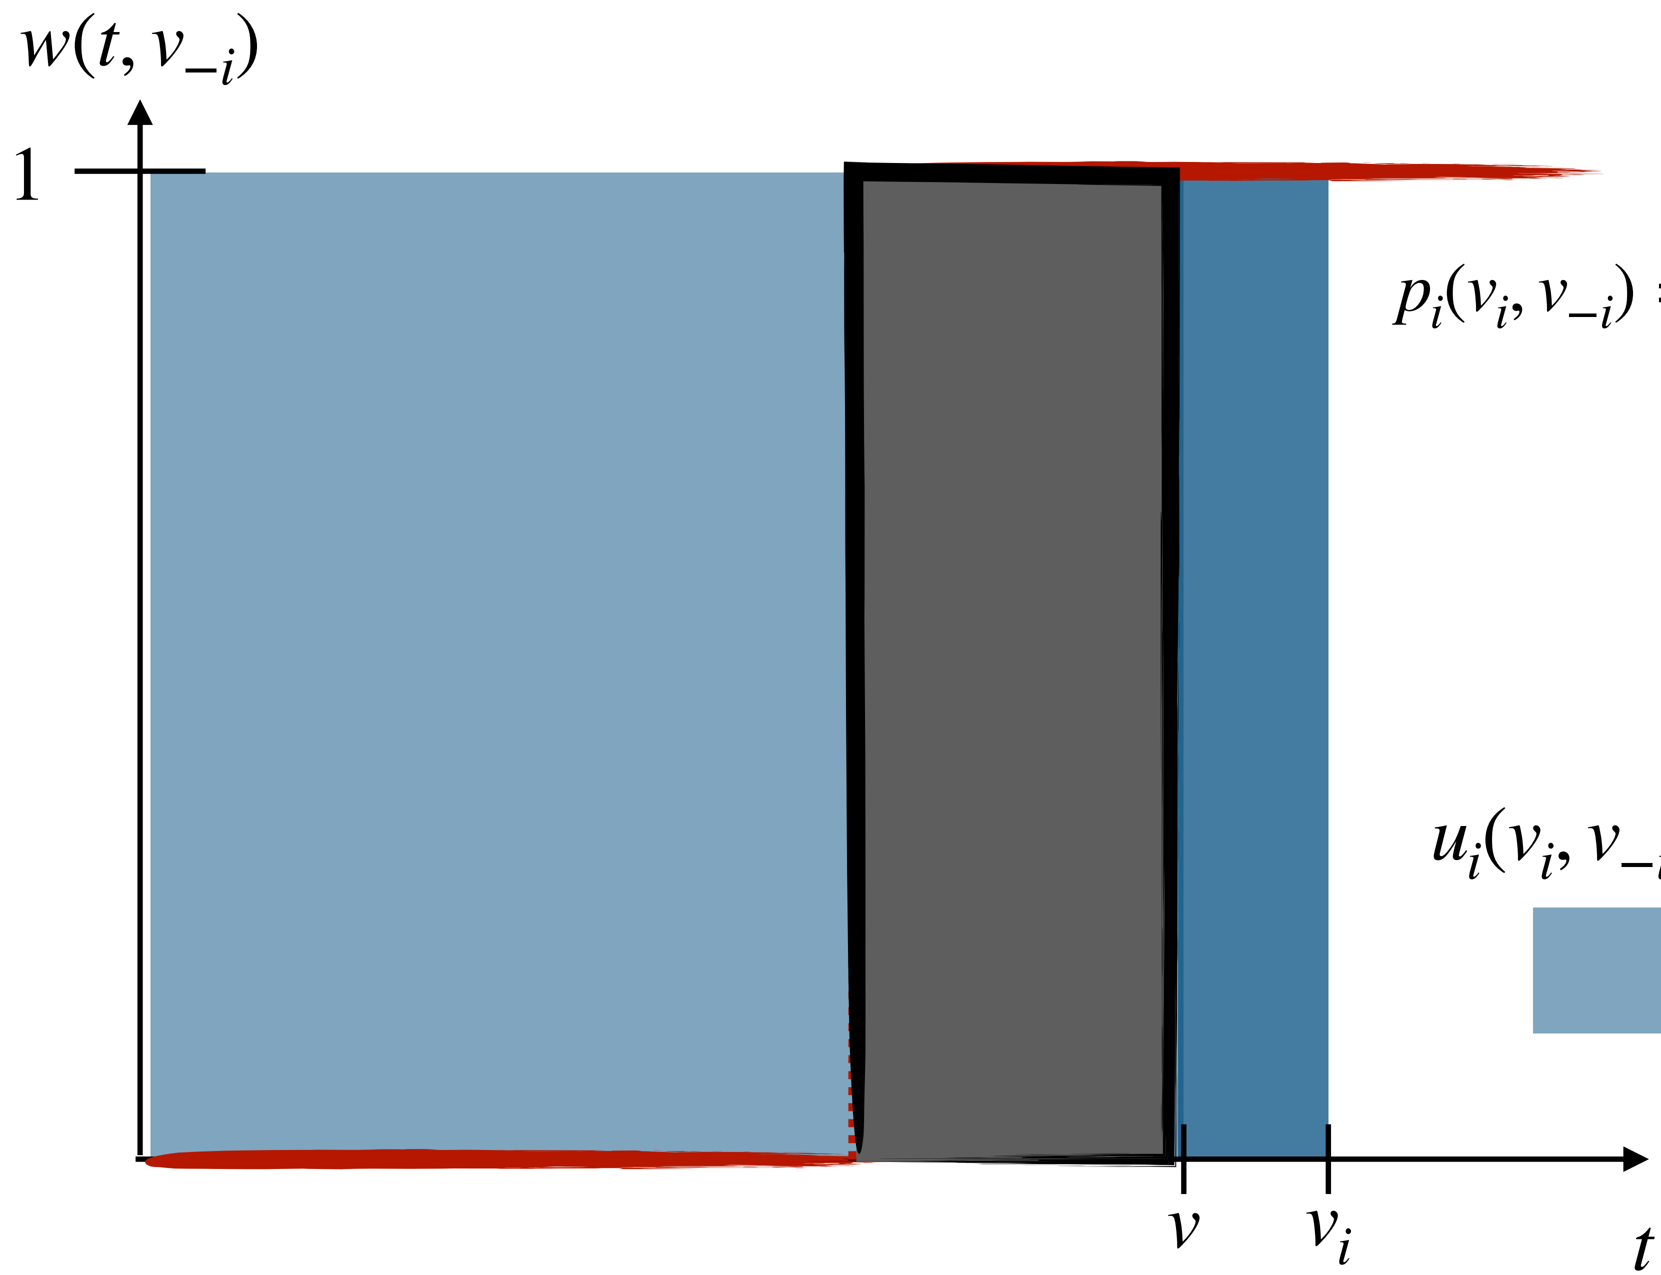

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

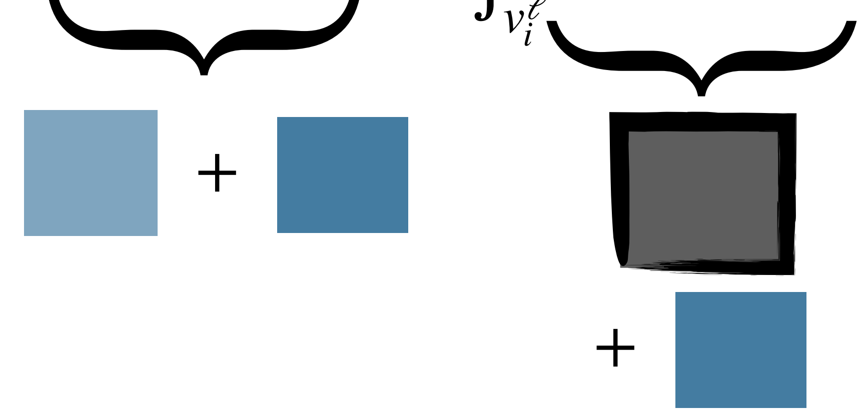

$$u_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

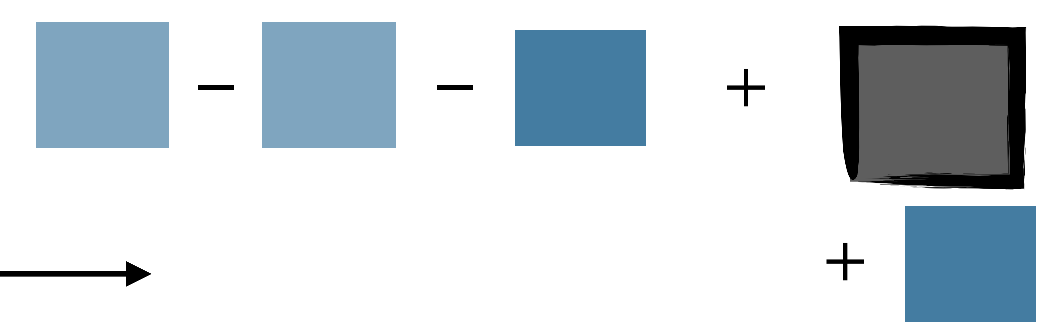

Some pictures

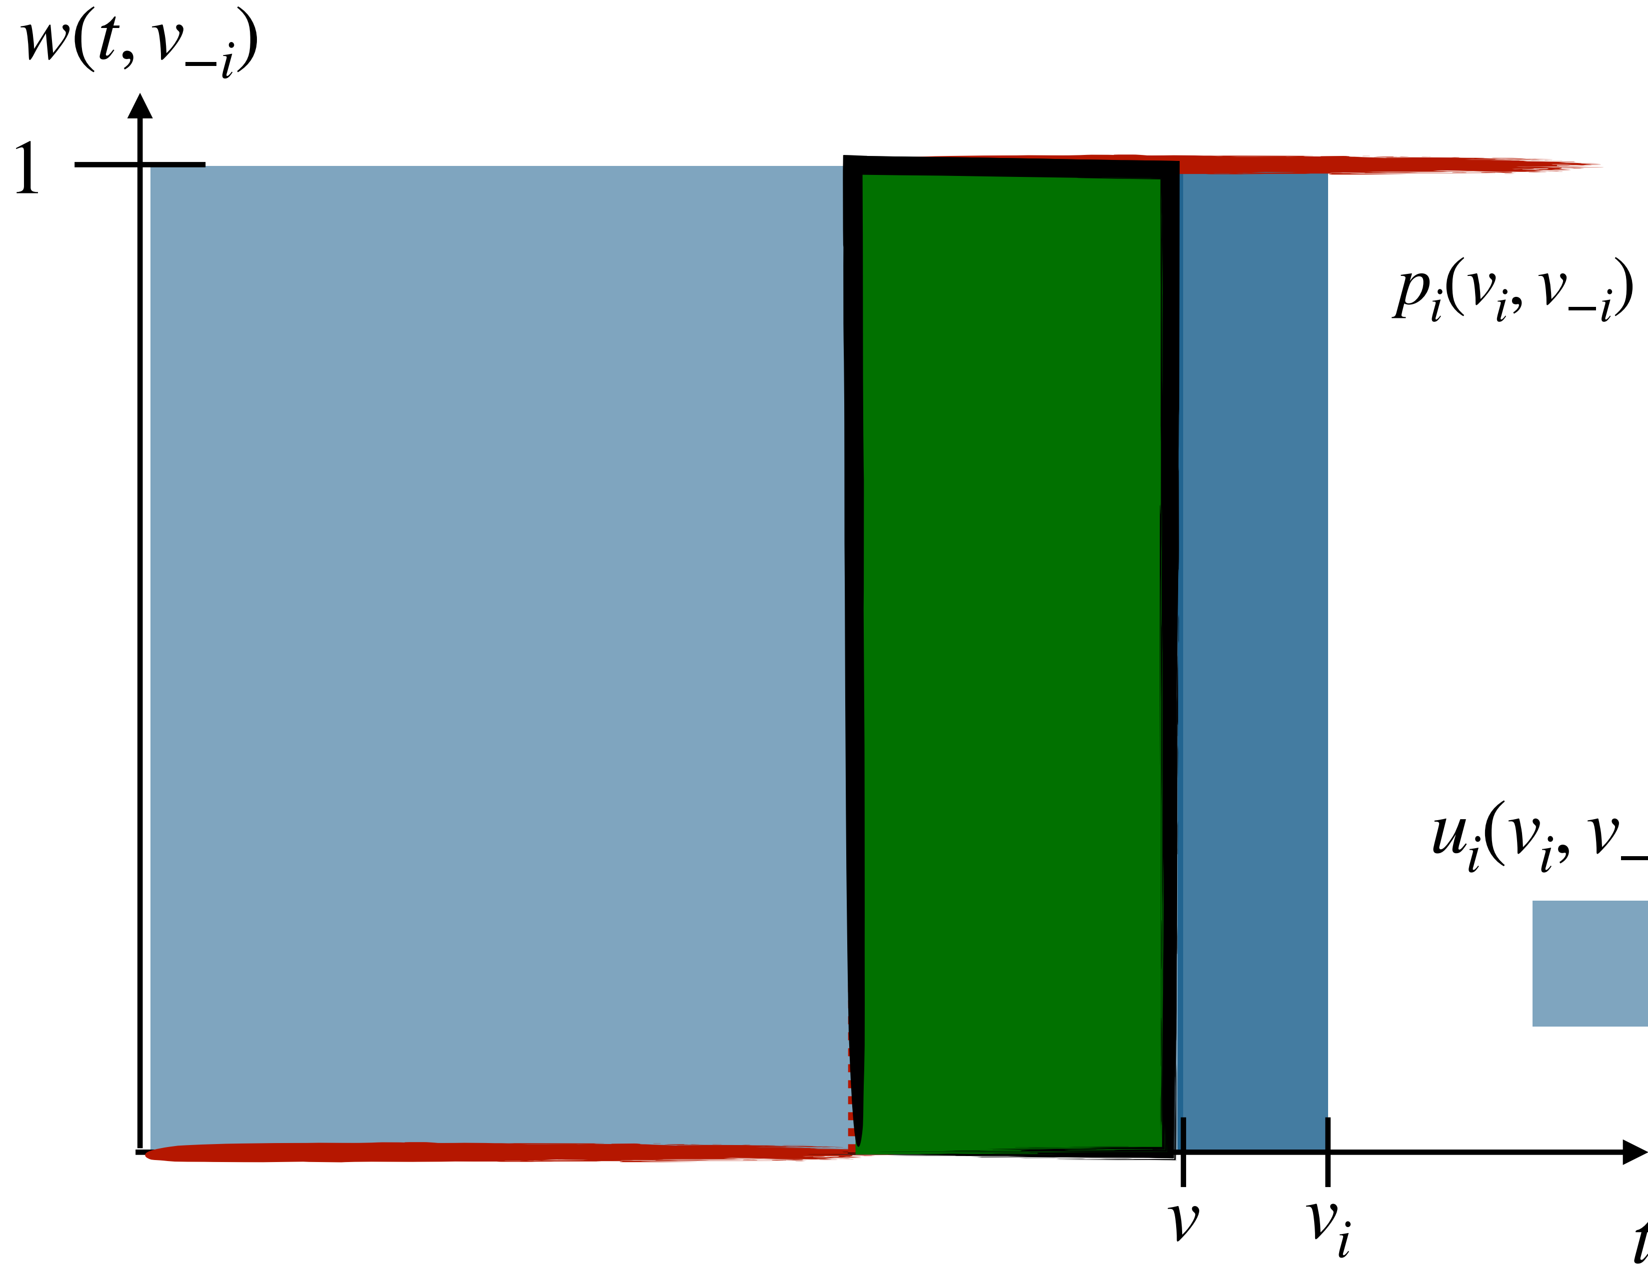

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

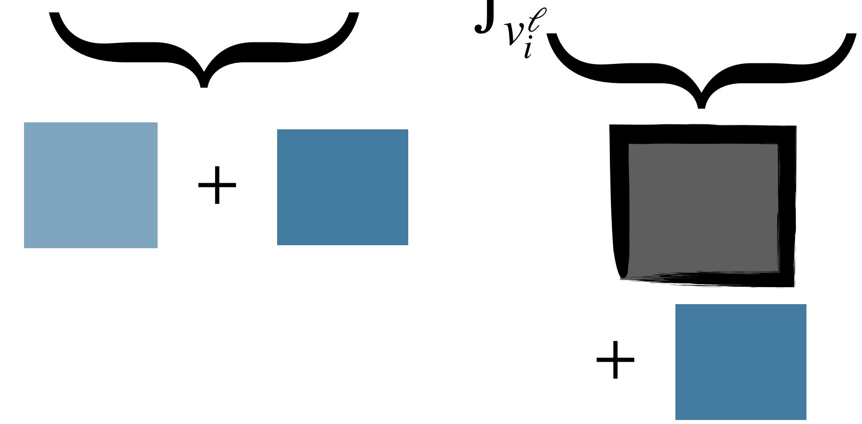

$$u_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

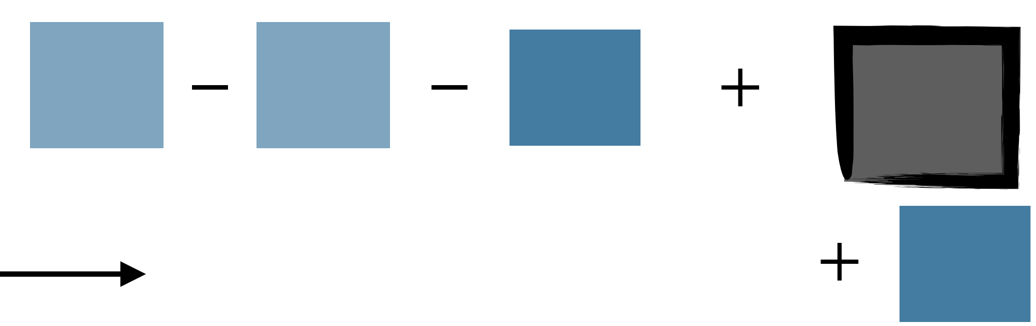

Some pictures

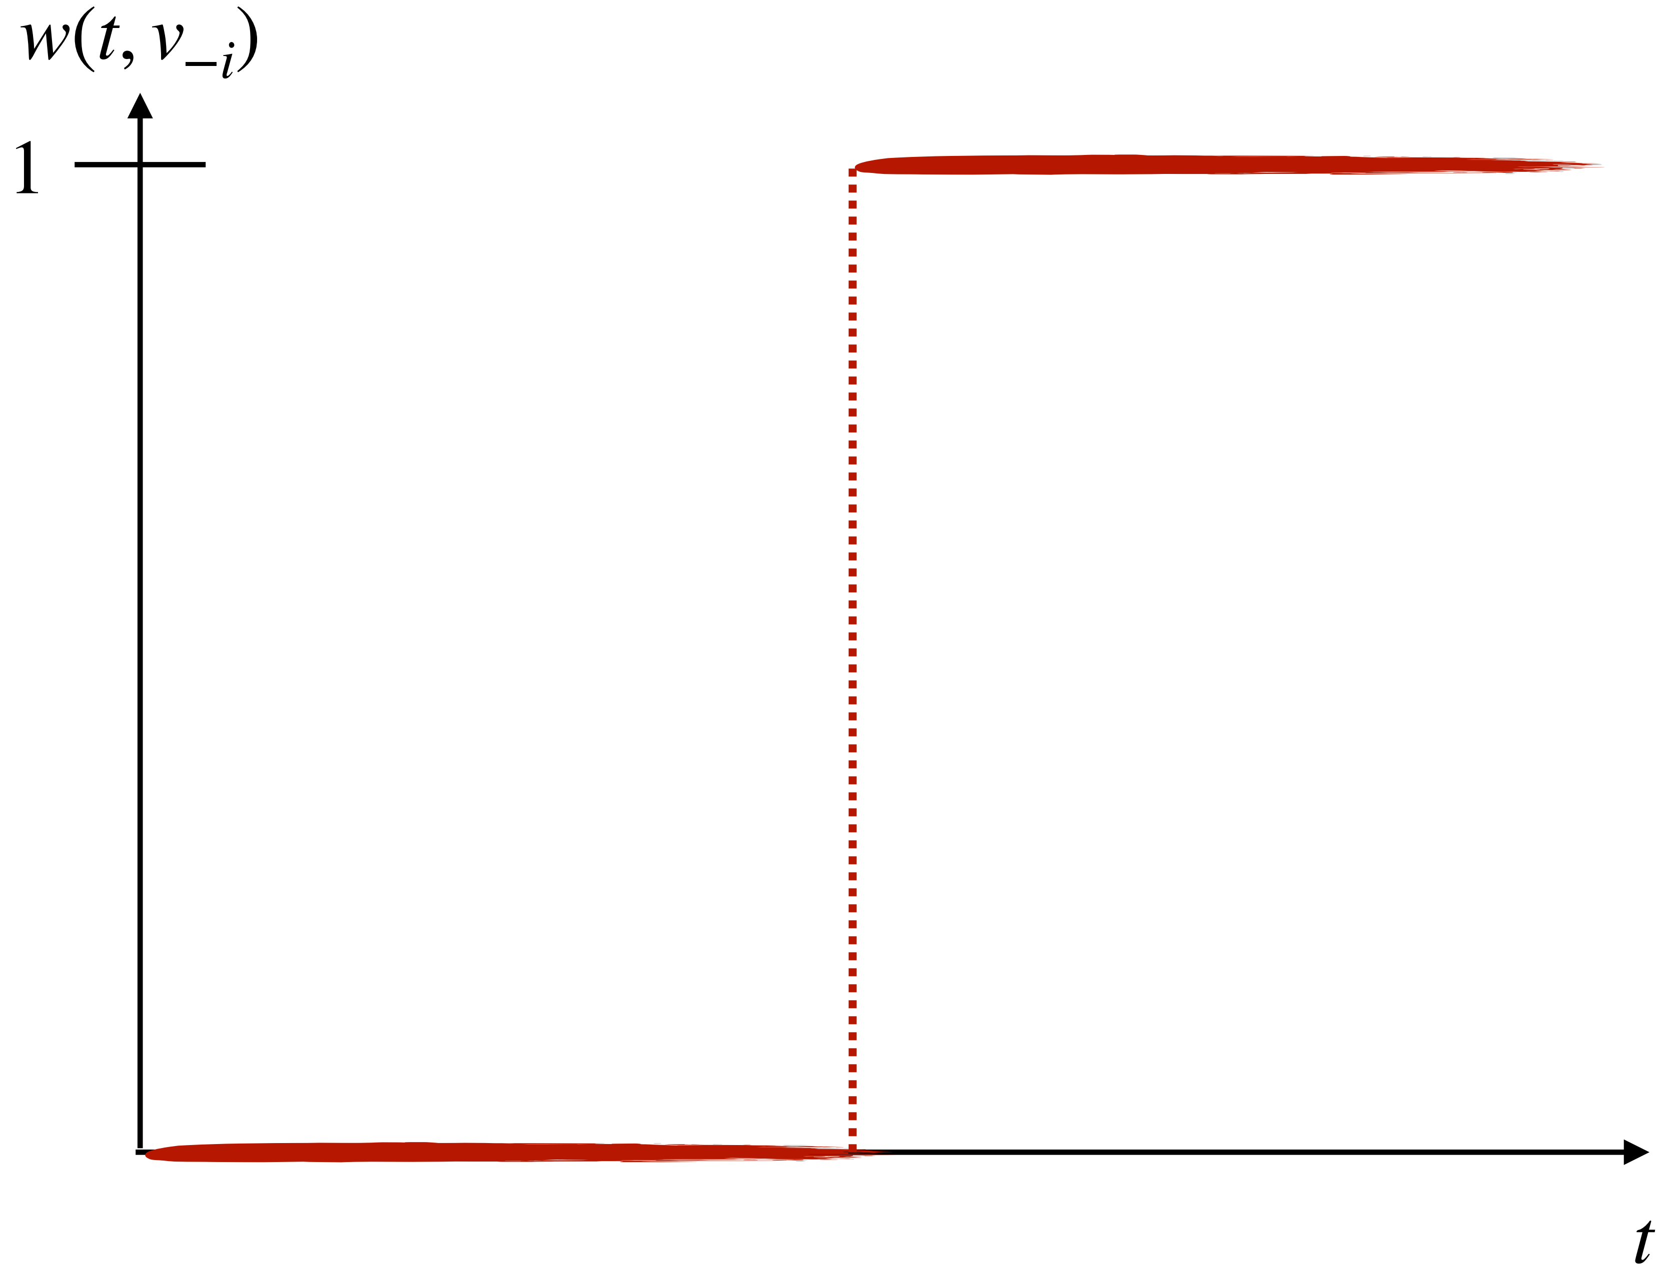

Some pictures

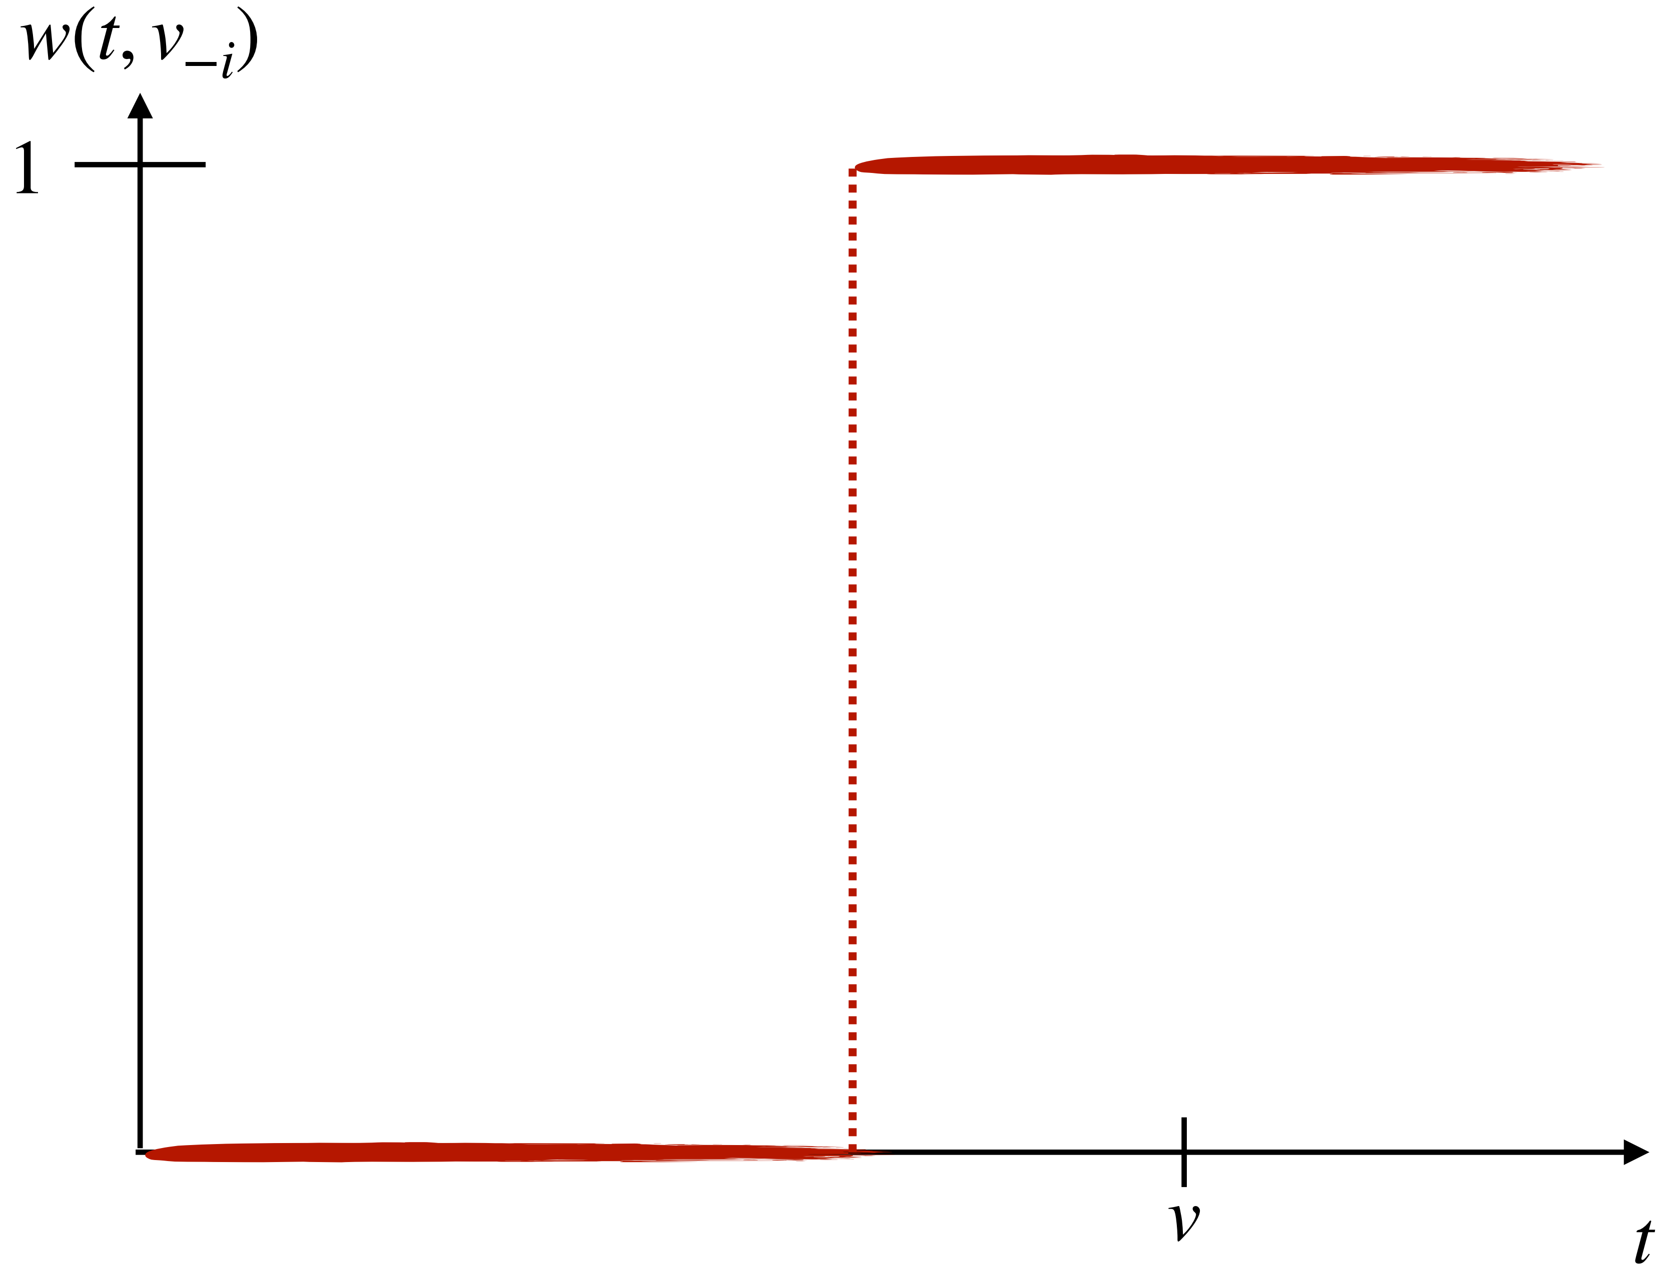

Some pictures

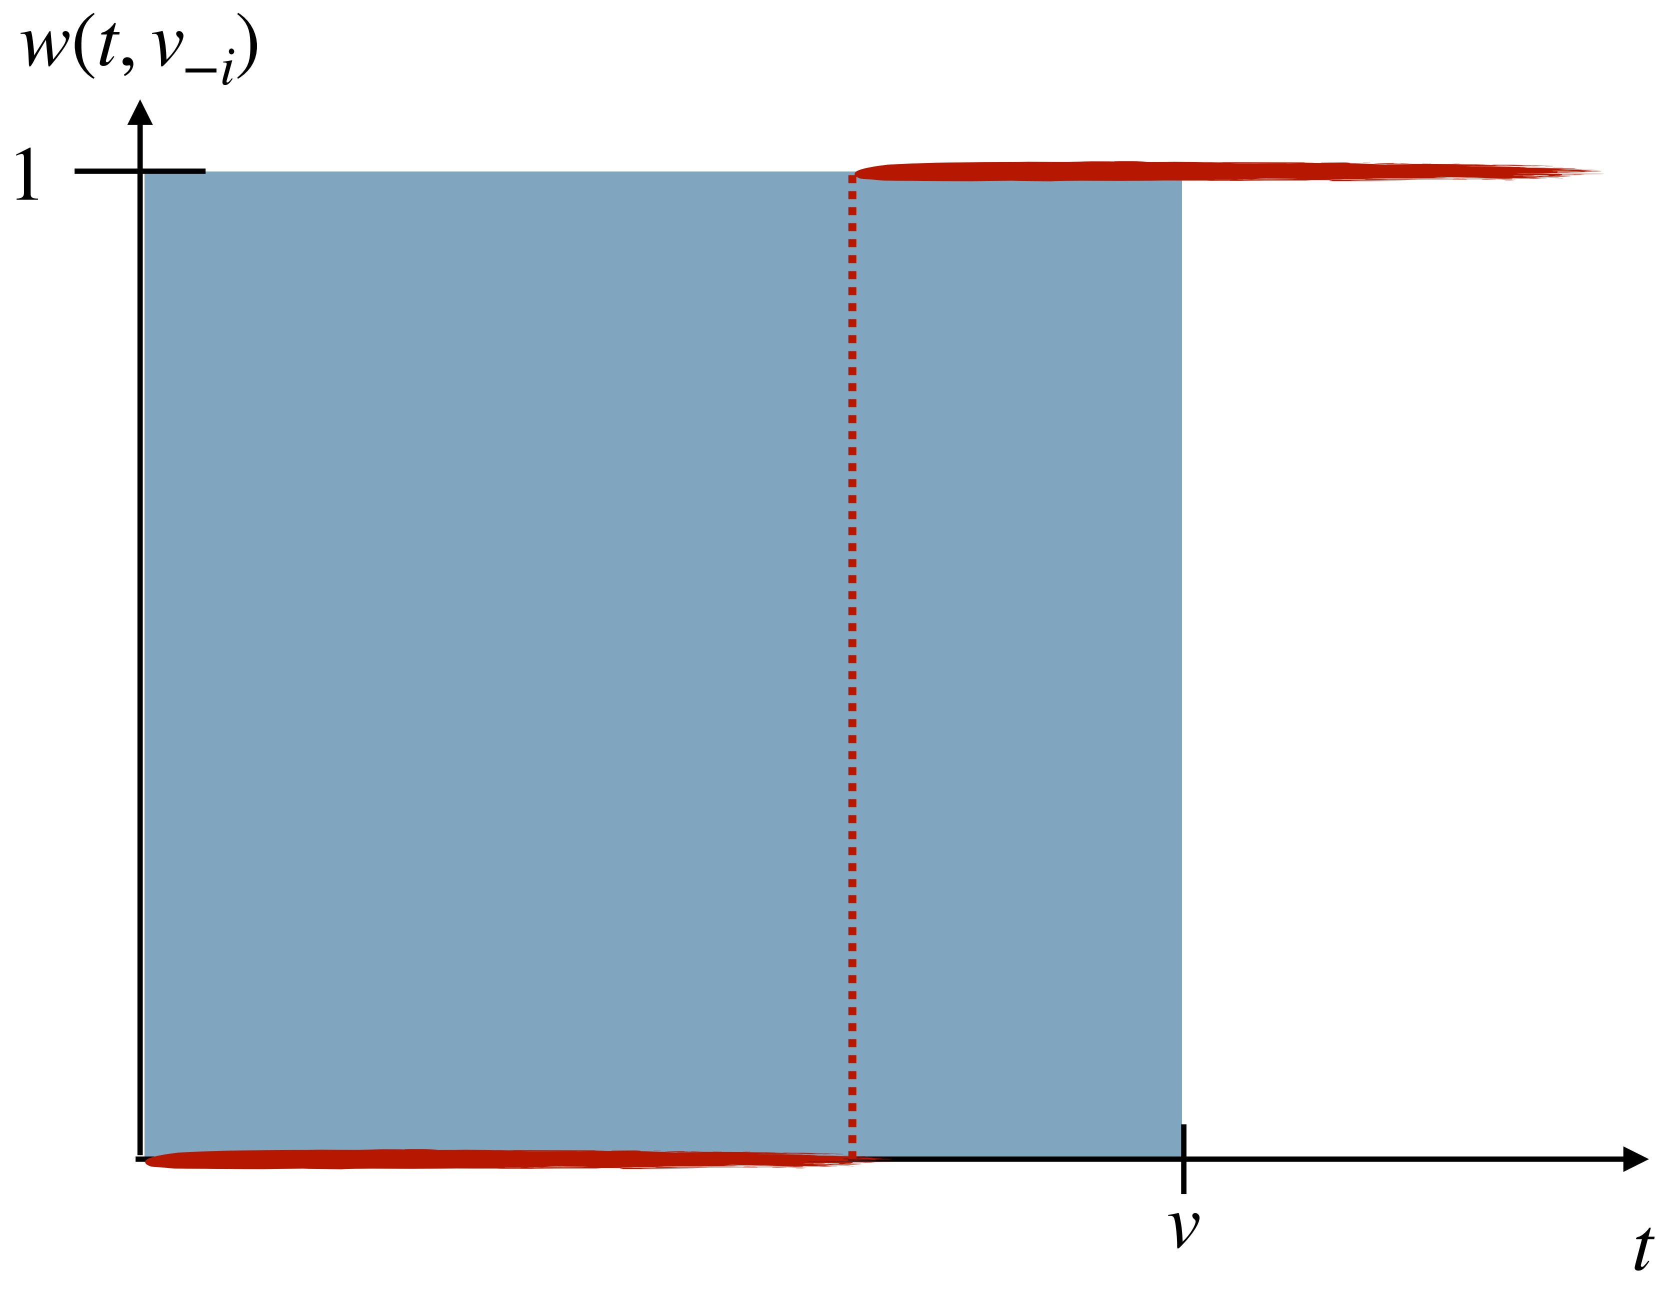

Some pictures

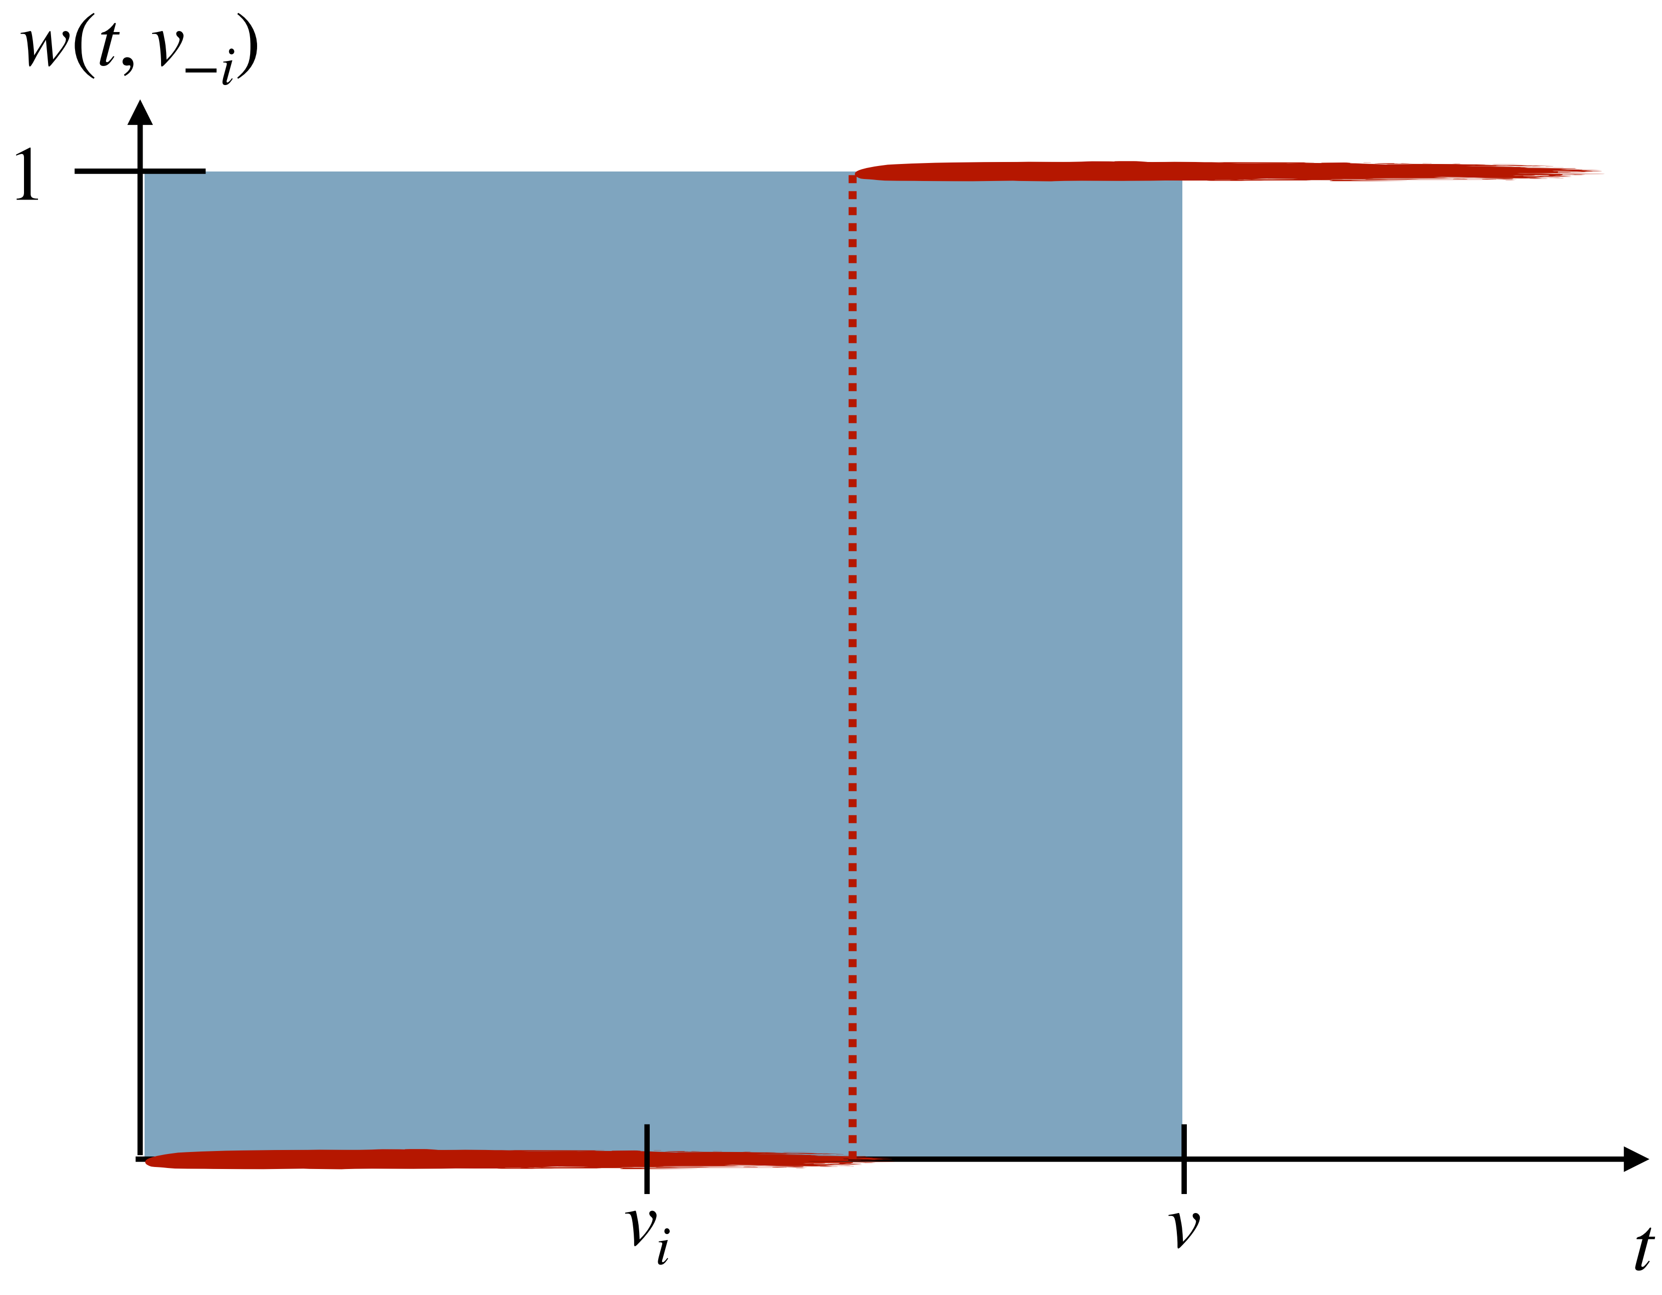

Some pictures

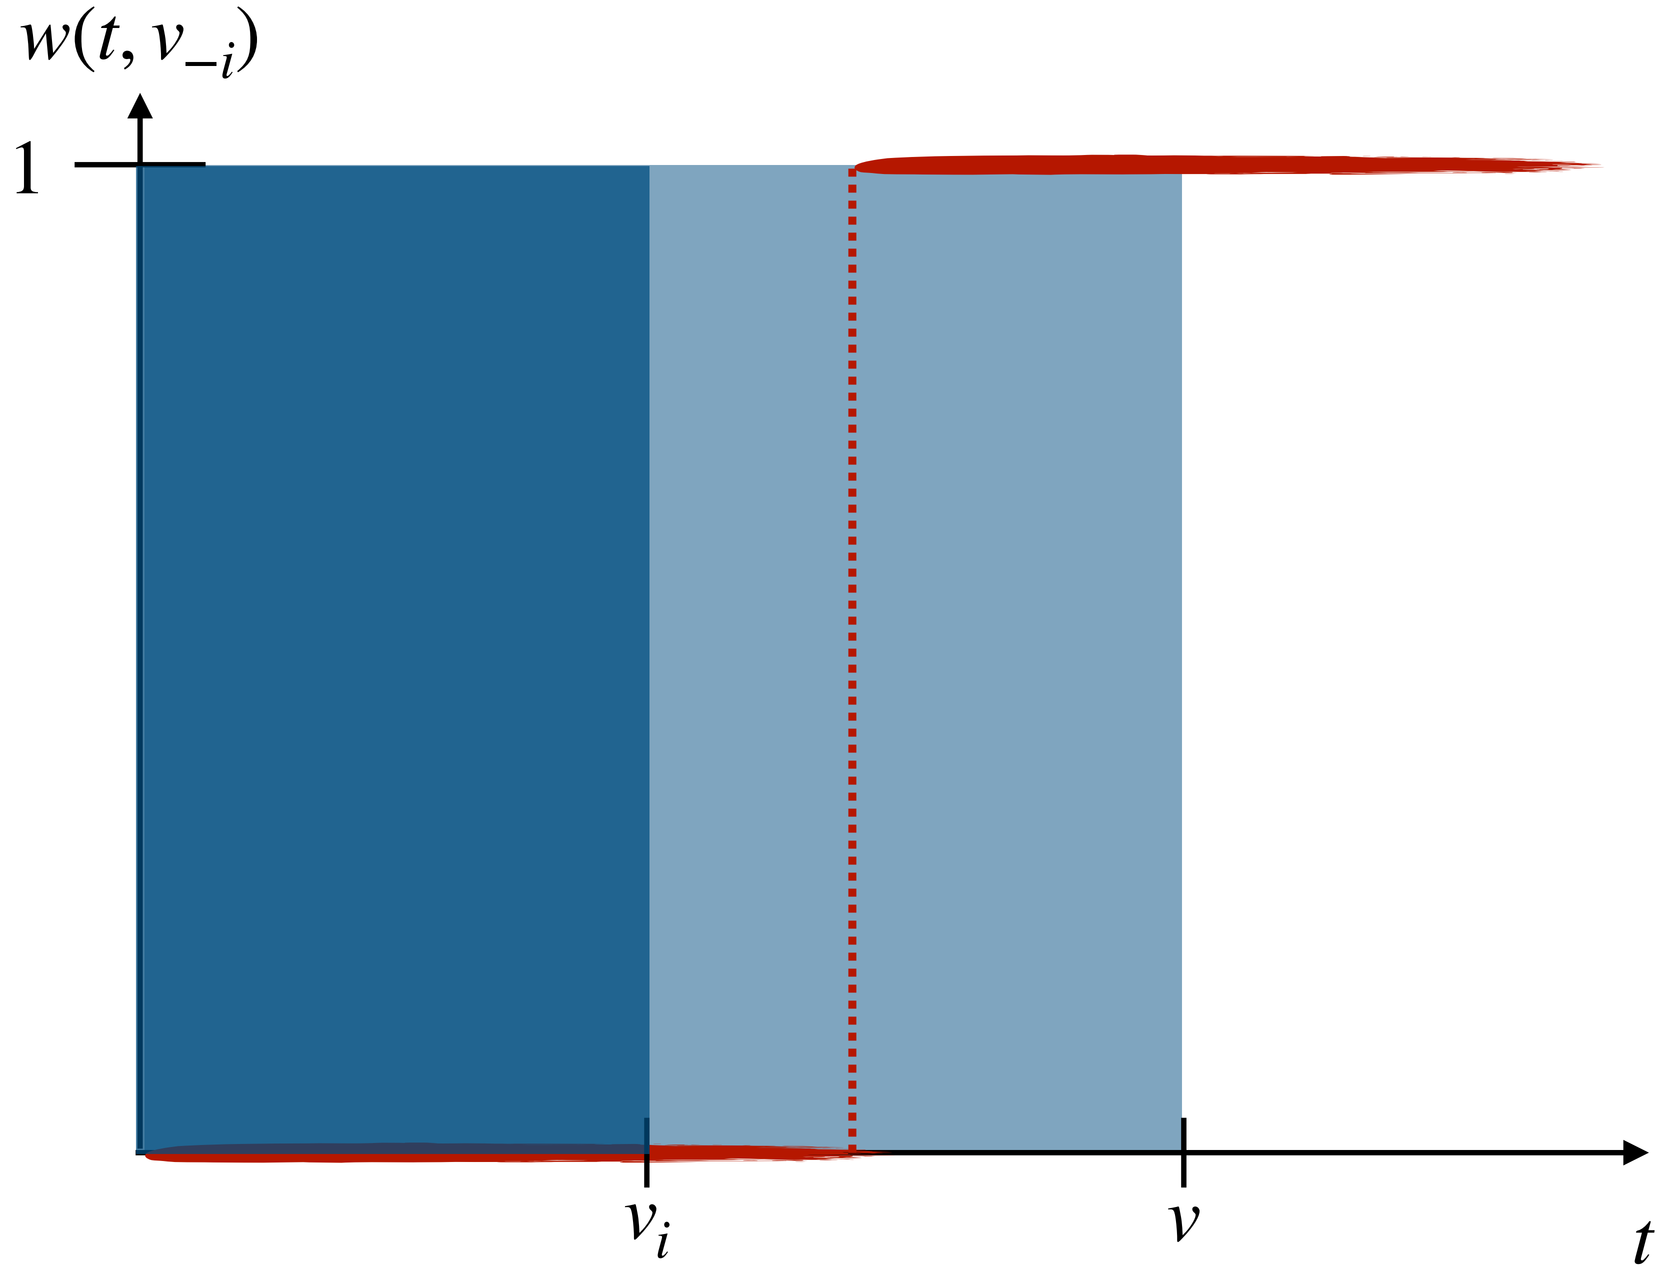

Some pictures

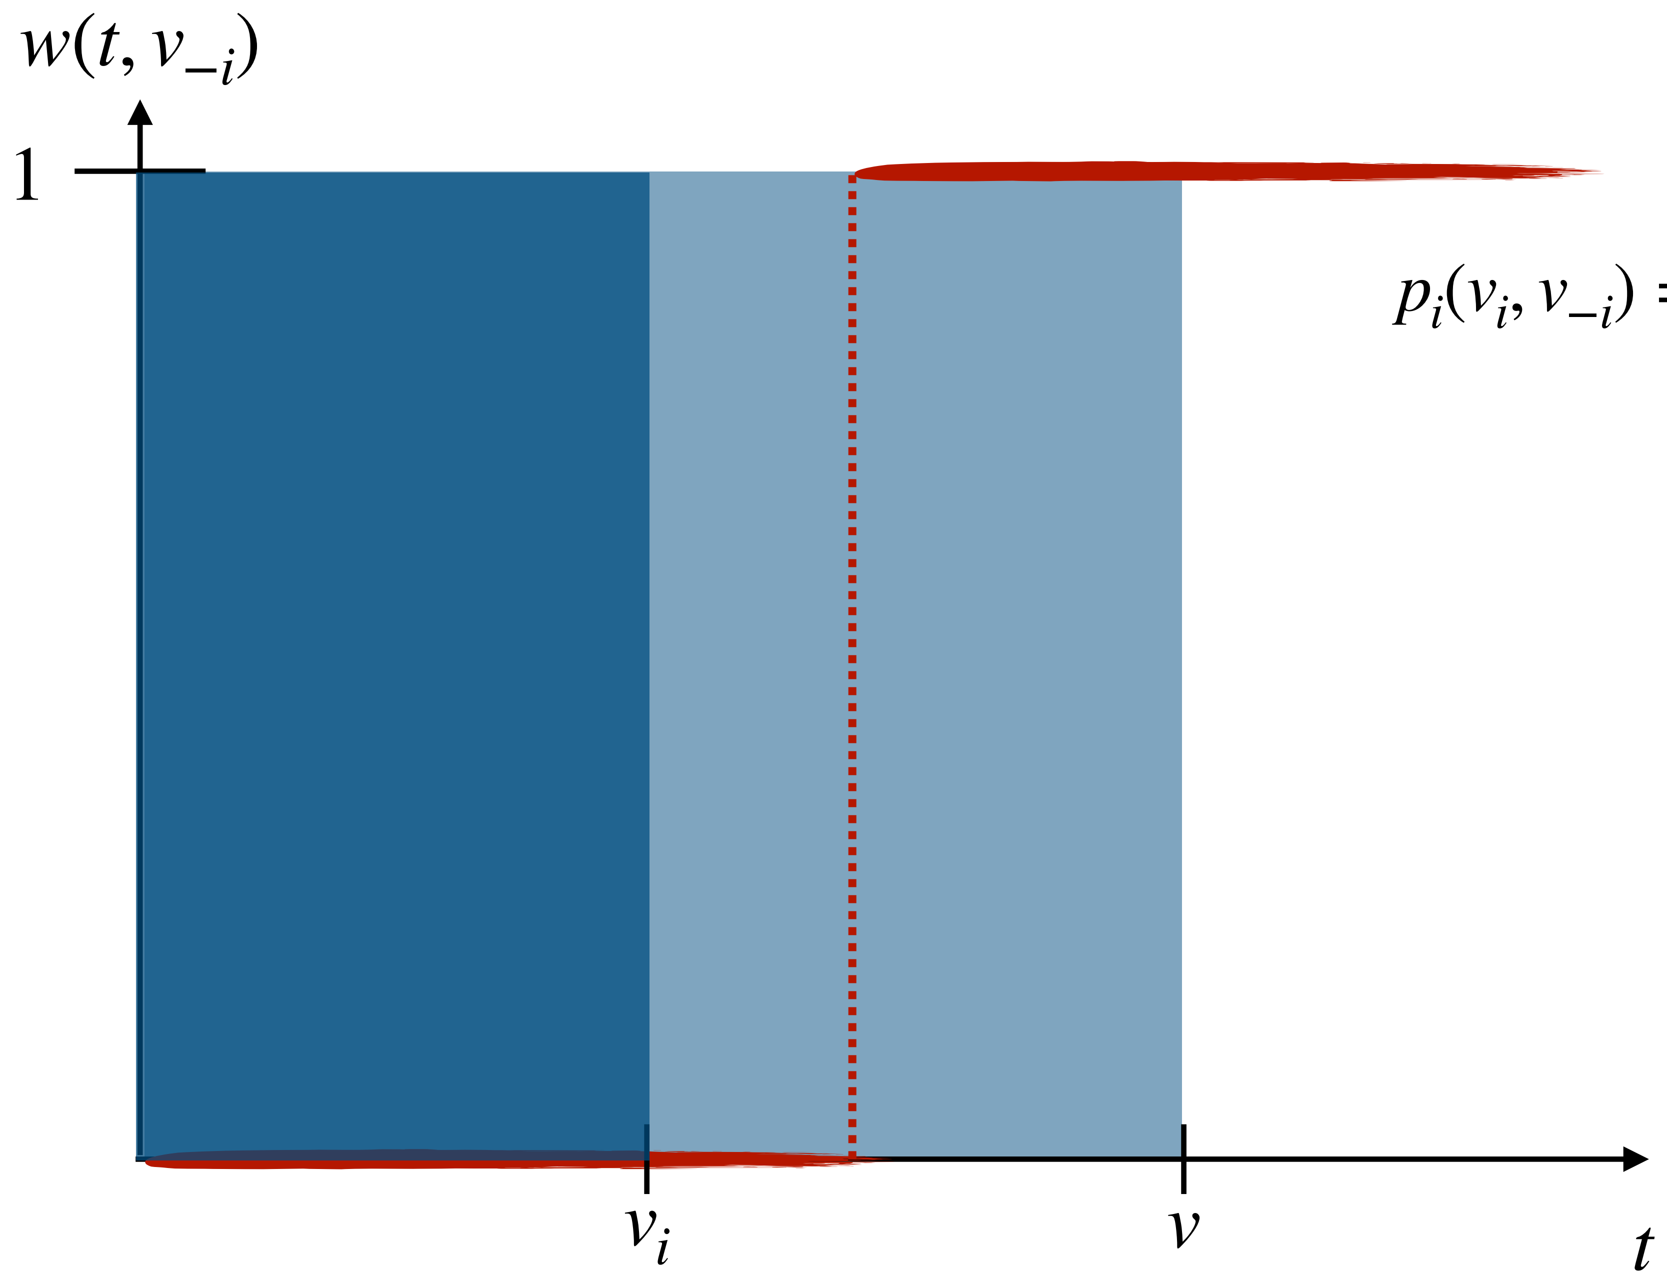

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

Some pictures

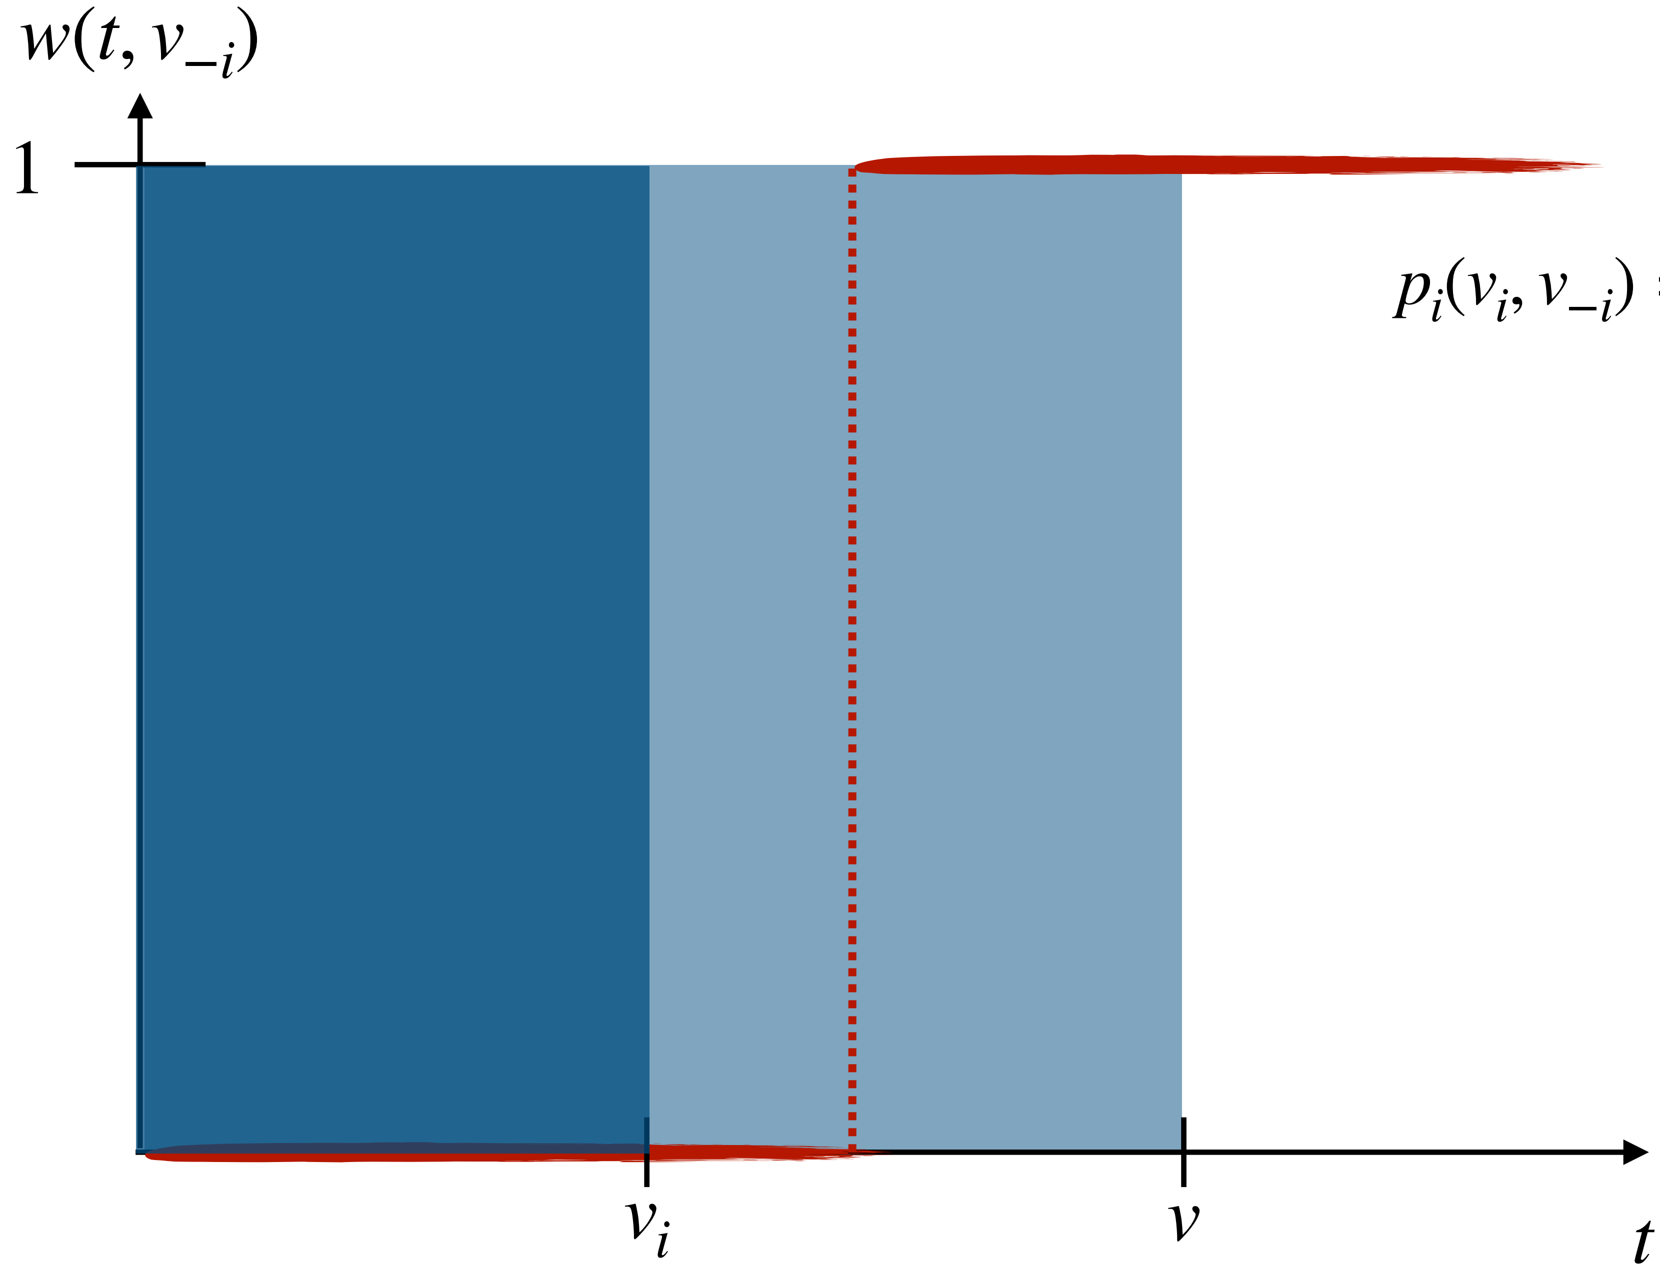

$$p_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})}_{\blacksquare} - \int_{v_i^\ell}^{v_i} \underbrace{w(t, v_{-i})dt}_{0}$$

Some pictures

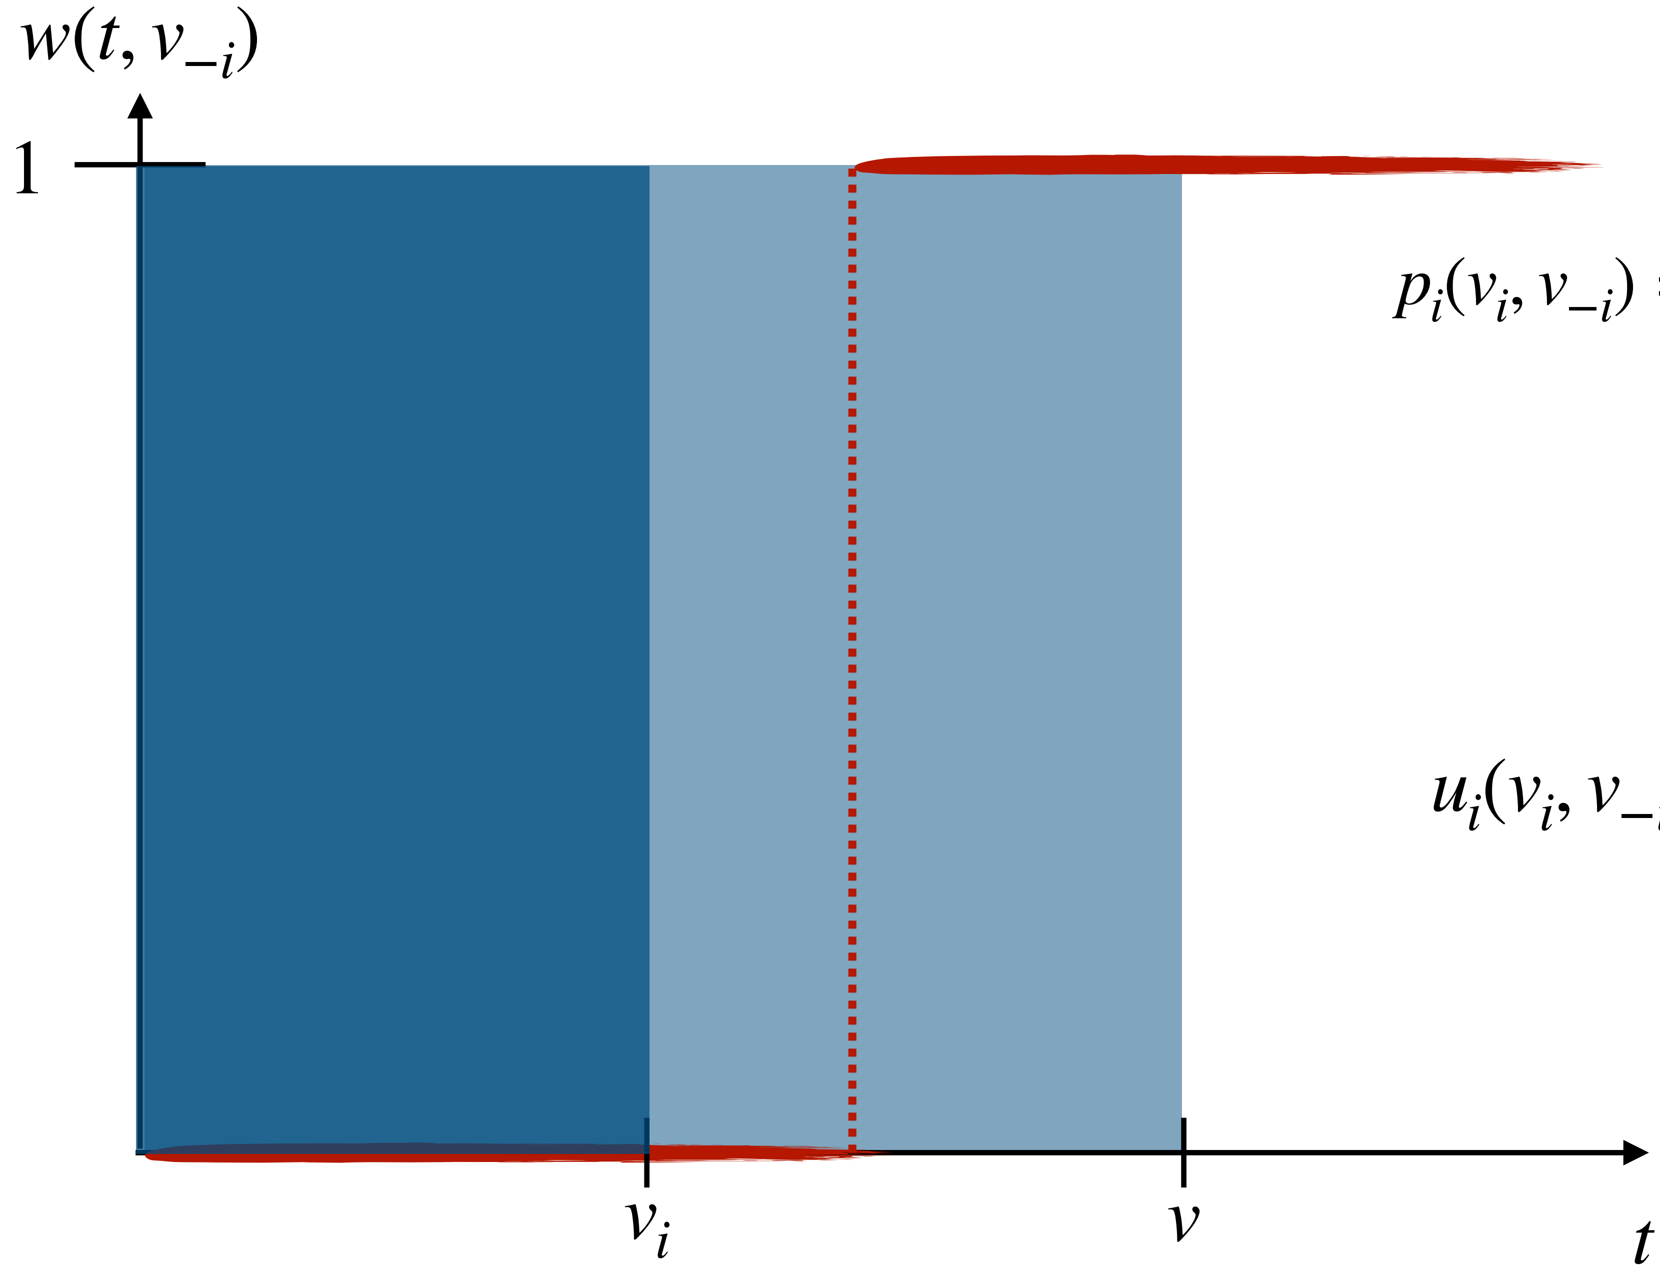

$$p_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})}_{\blacksquare} - \underbrace{\int_{v_i^\ell}^{v_i} w(t, v_{-i})dt}_{0}$$

$$u_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

Some pictures

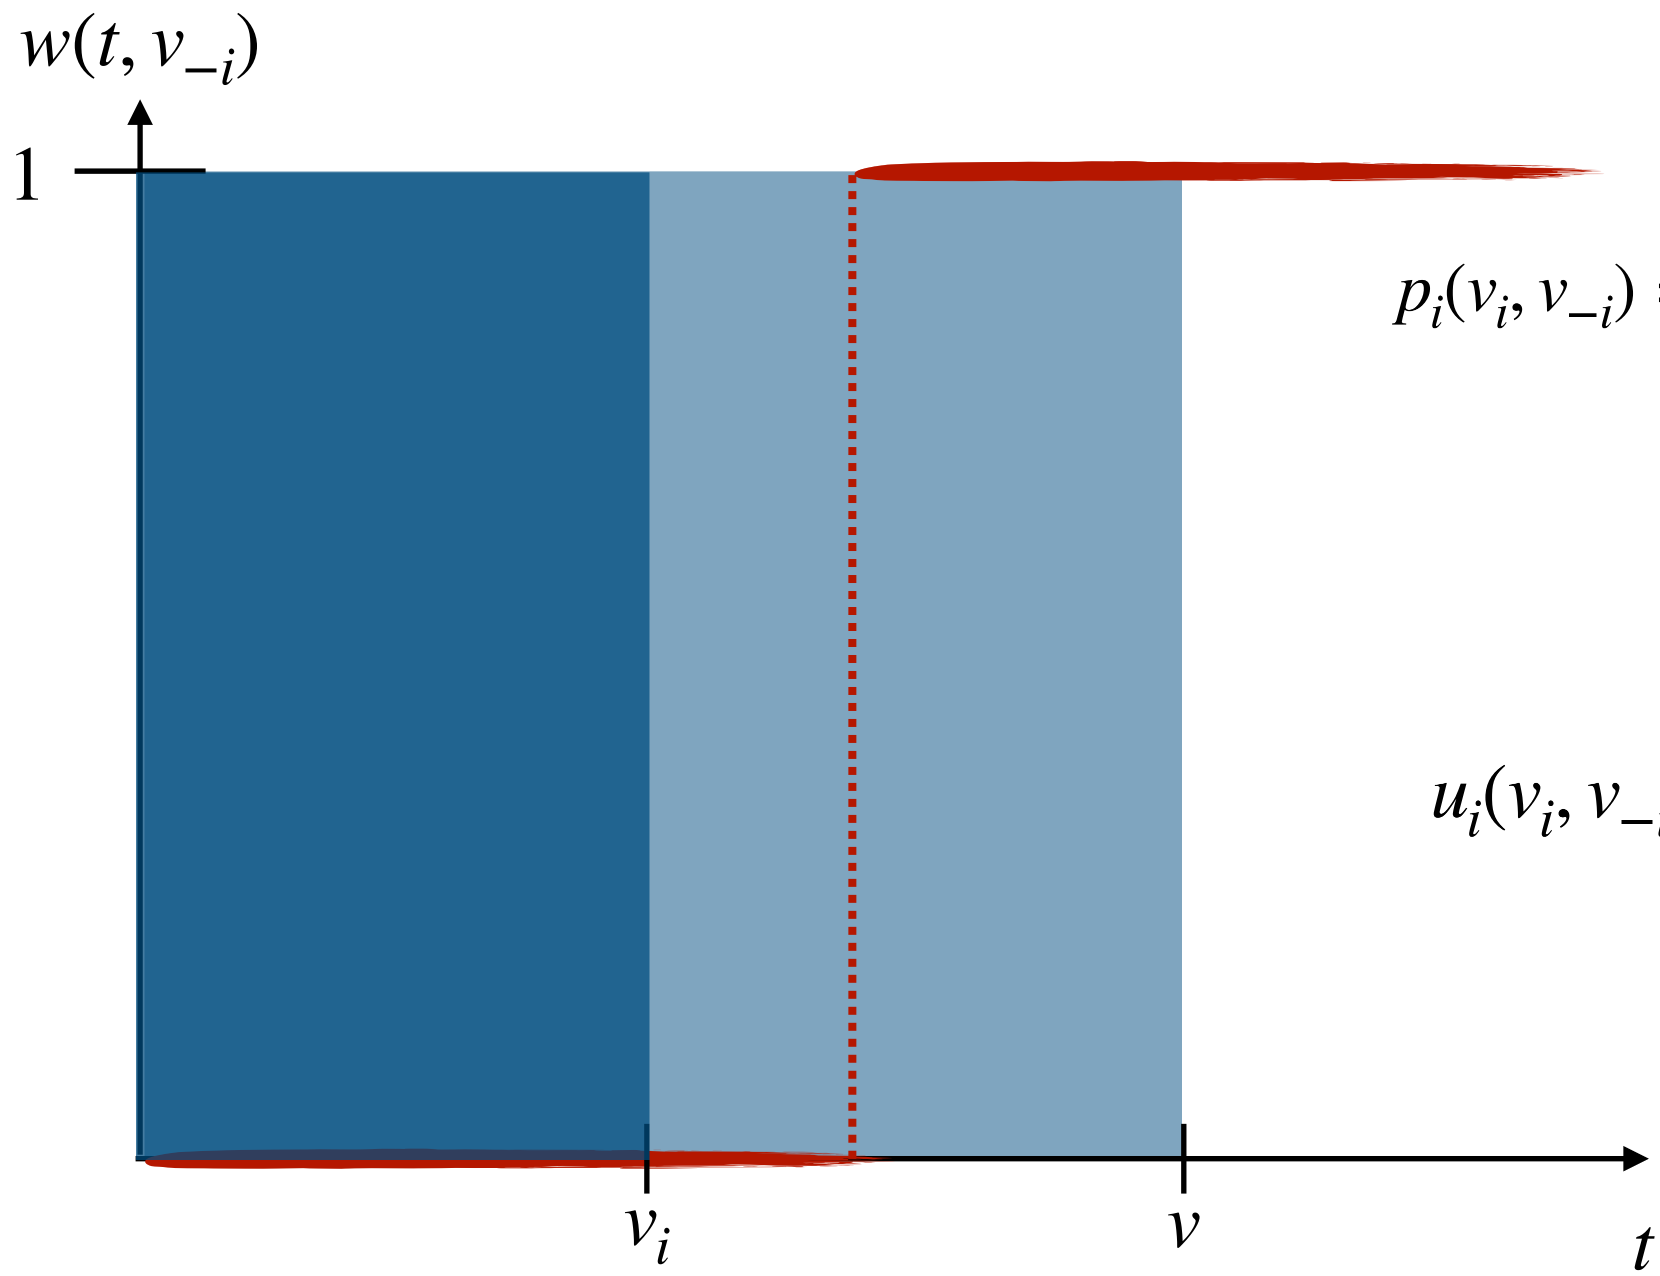

$$p_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})}_{\blacksquare} - \int_{v_i^\ell}^{v_i} \underbrace{w(t, v_{-i})dt}_{0}$$

$$u_i(v_i, v_{-i}) = \underbrace{v_i \cdot w_i(v_i, v_{-i})}_{\blacksquare} - p_i(v_i, v_{-i})$$

Some pictures

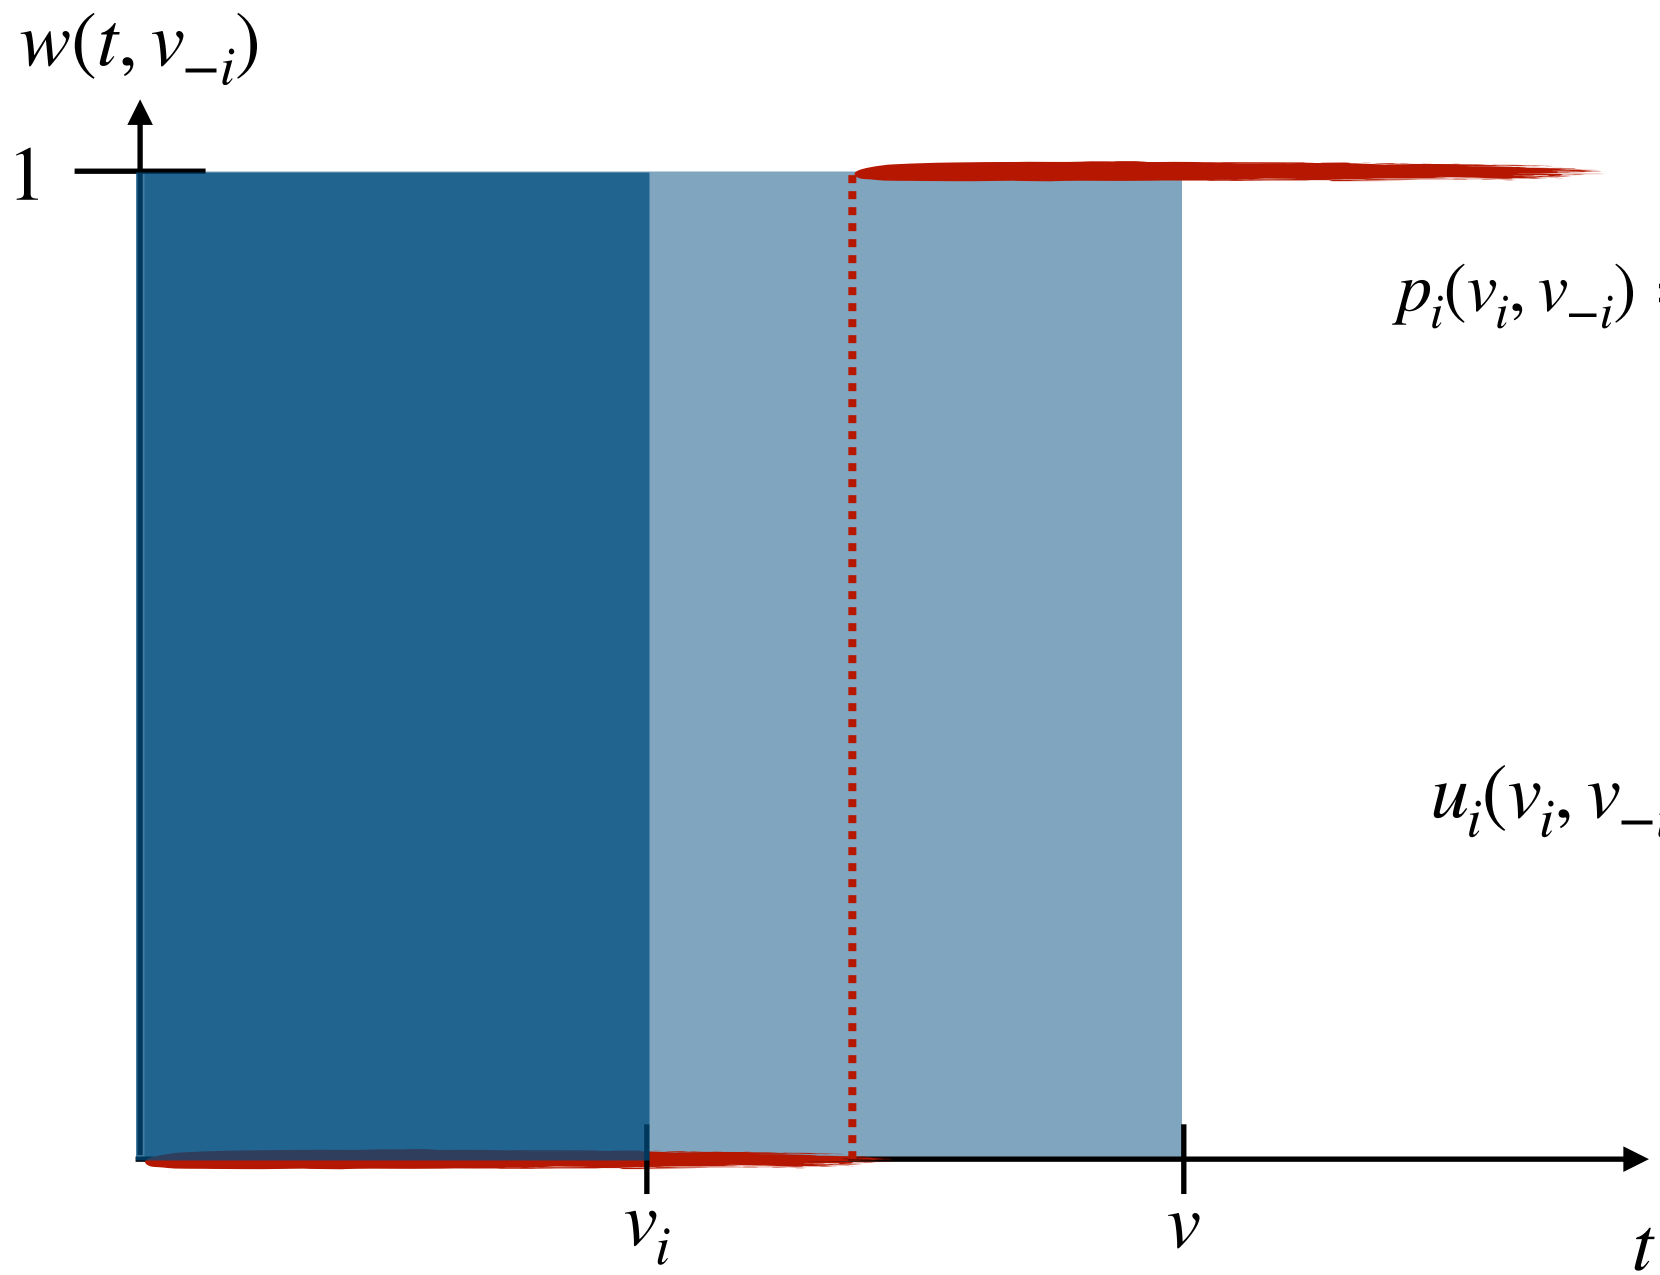

$$p_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^\ell}^{v_i} w(t, v_{-i})dt$$

0

$$u_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

Some pictures

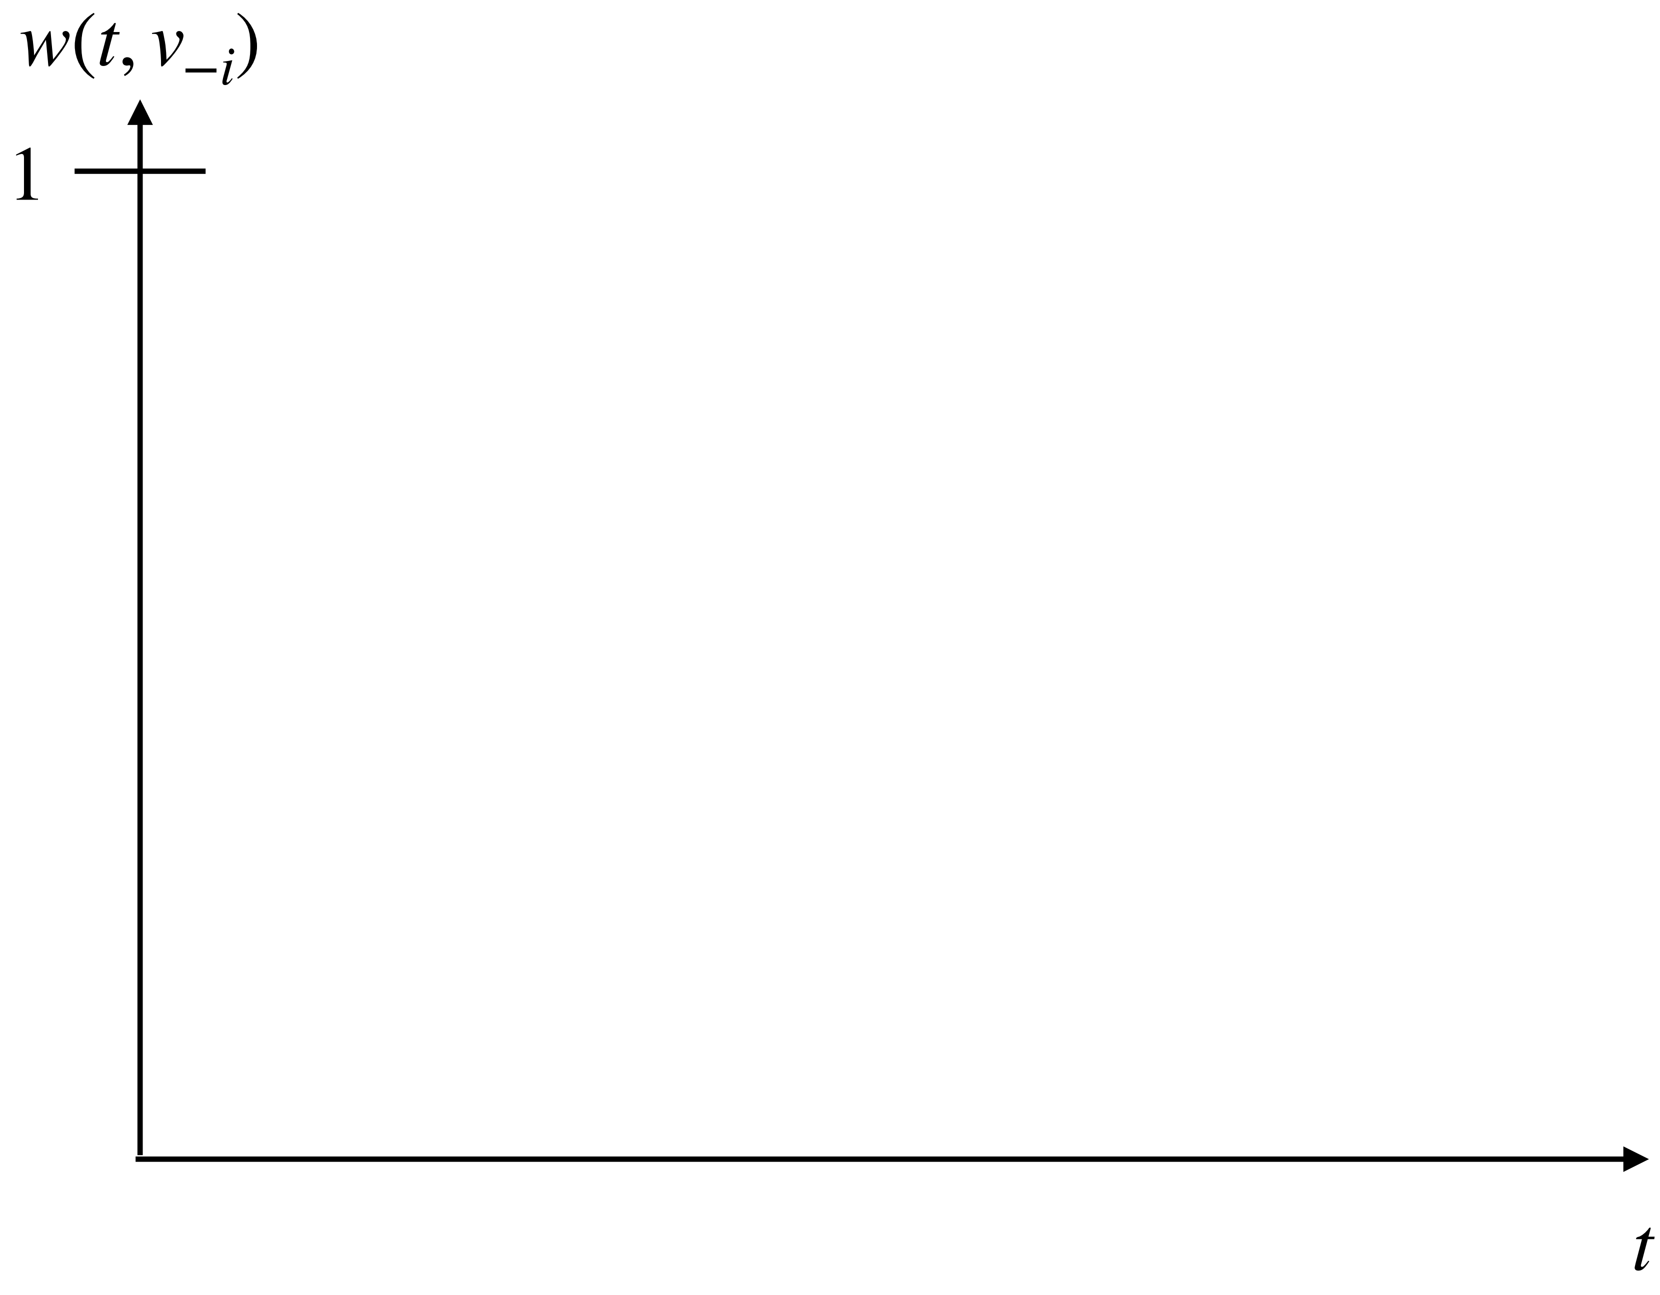

Some pictures

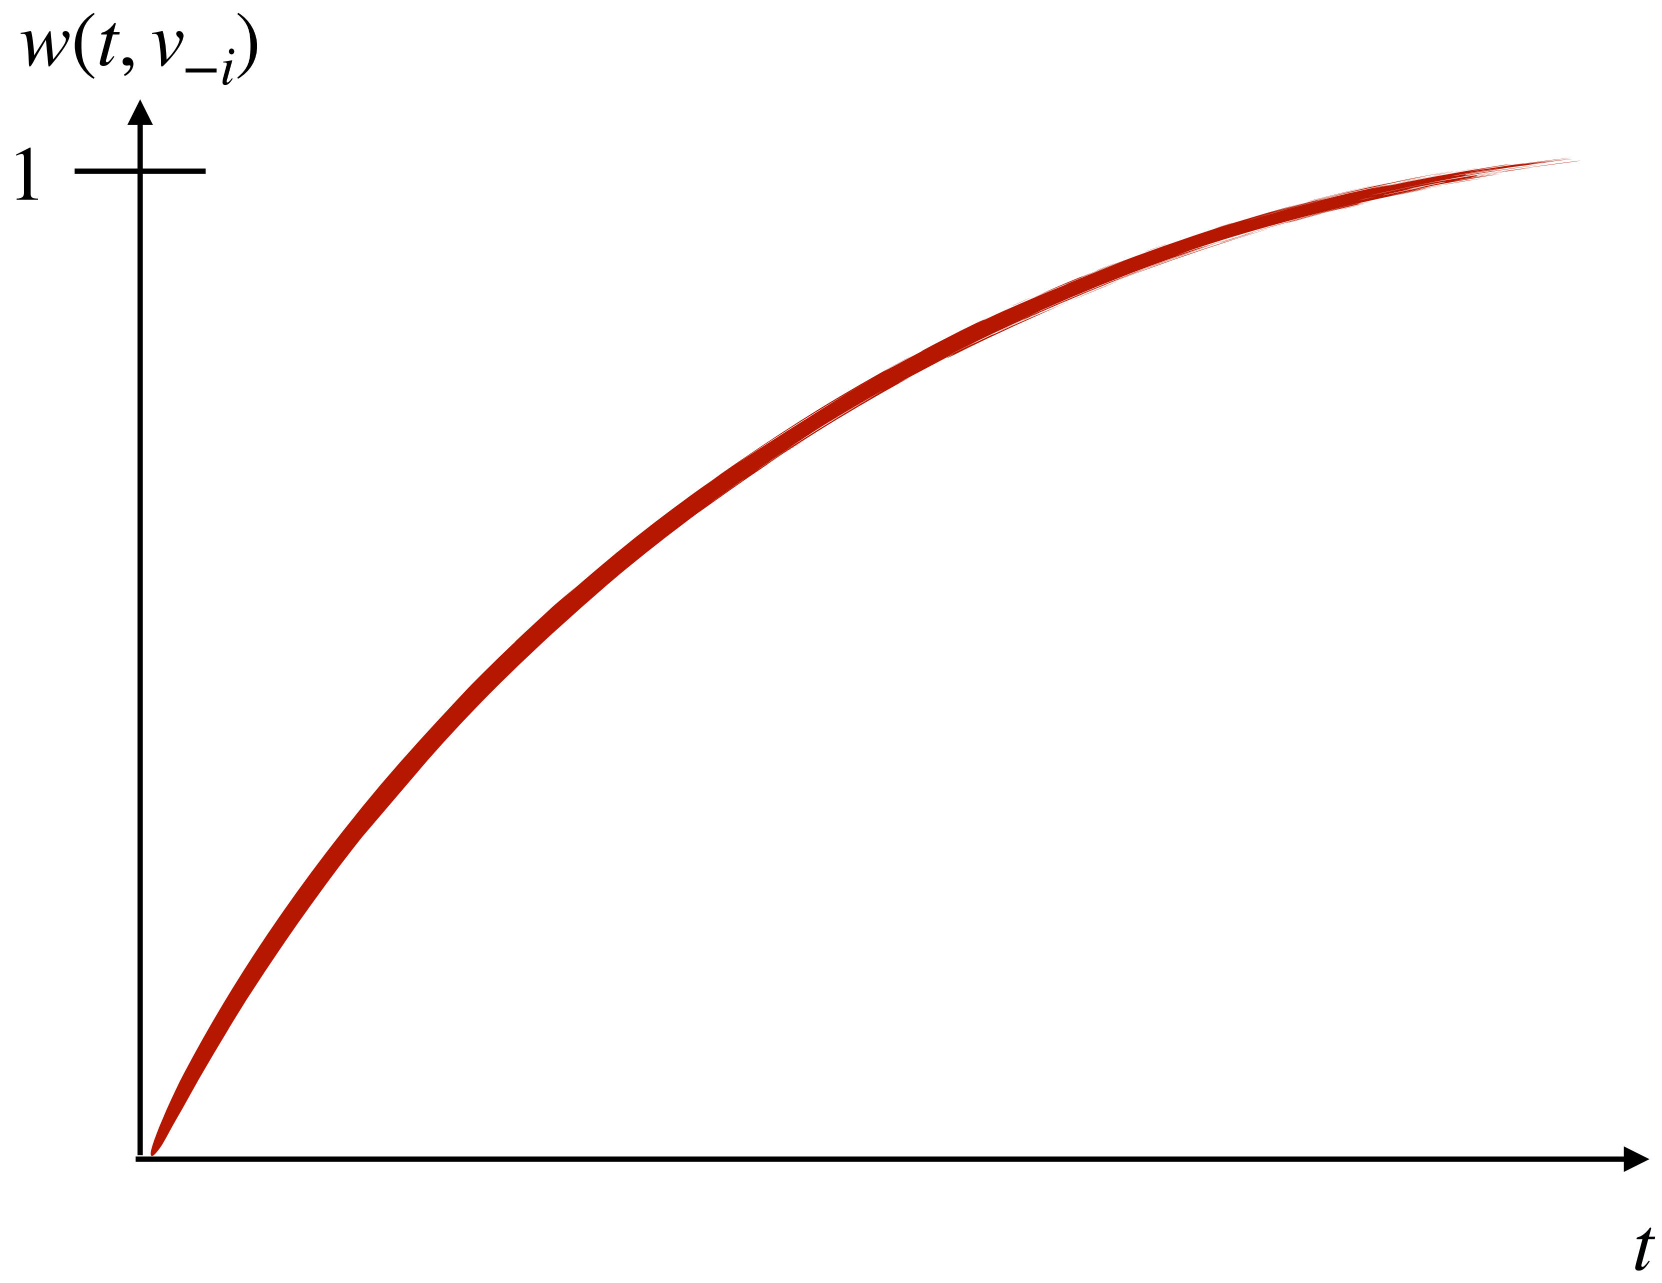

Some pictures

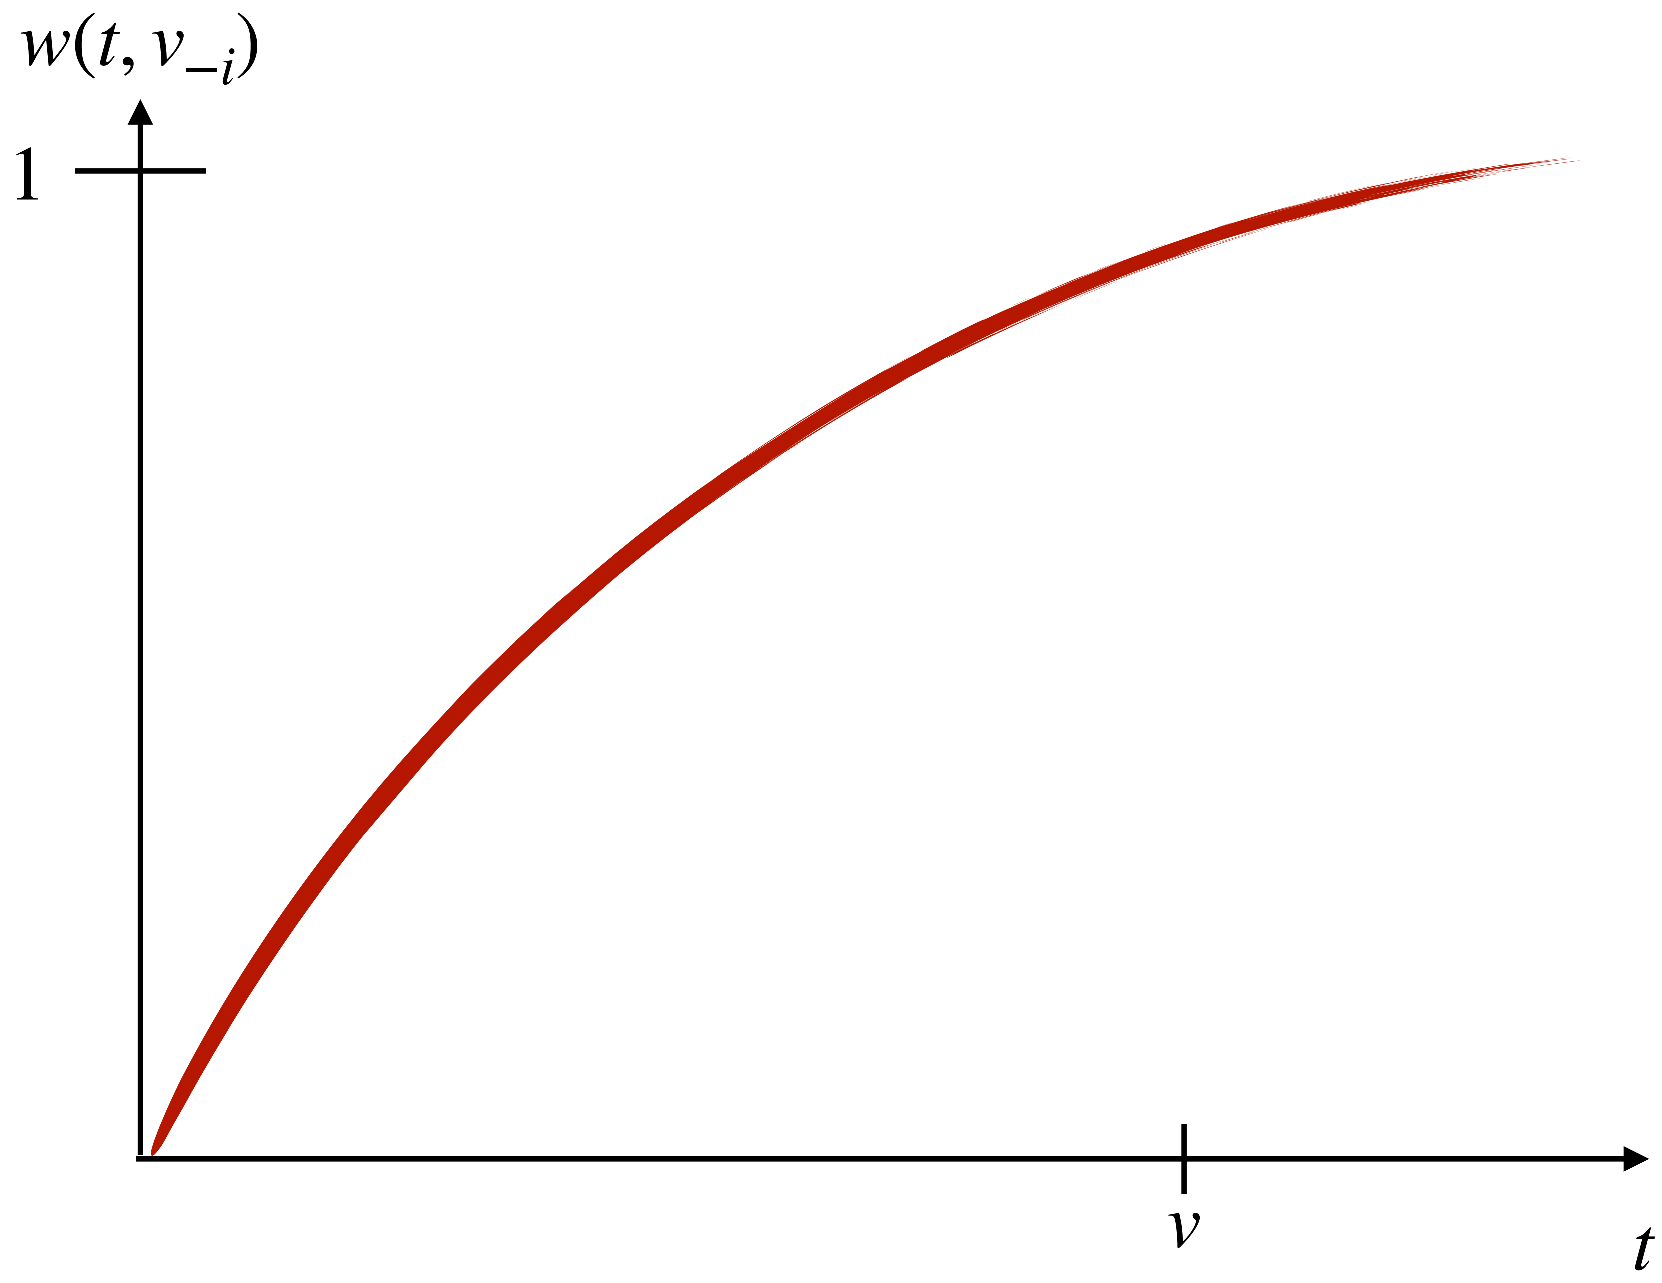

Some pictures

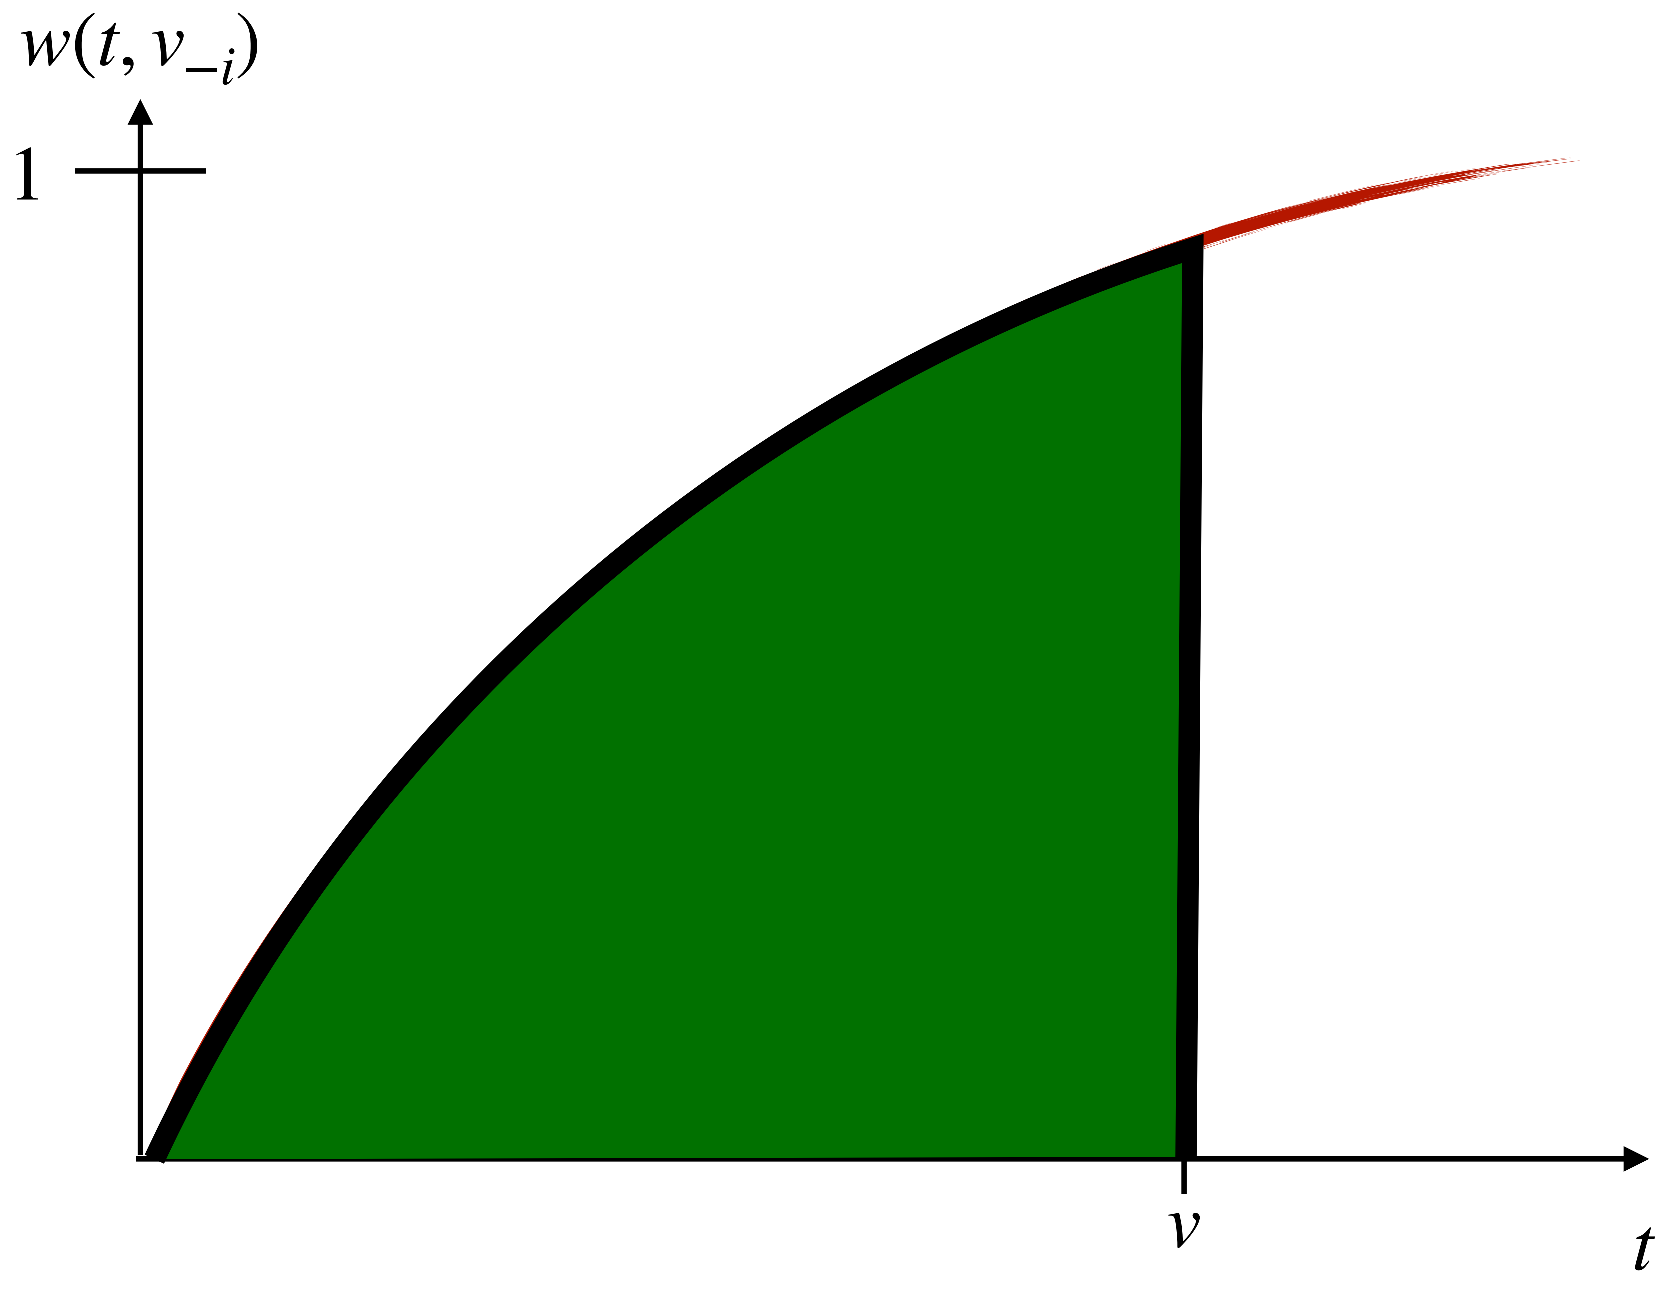

Some pictures

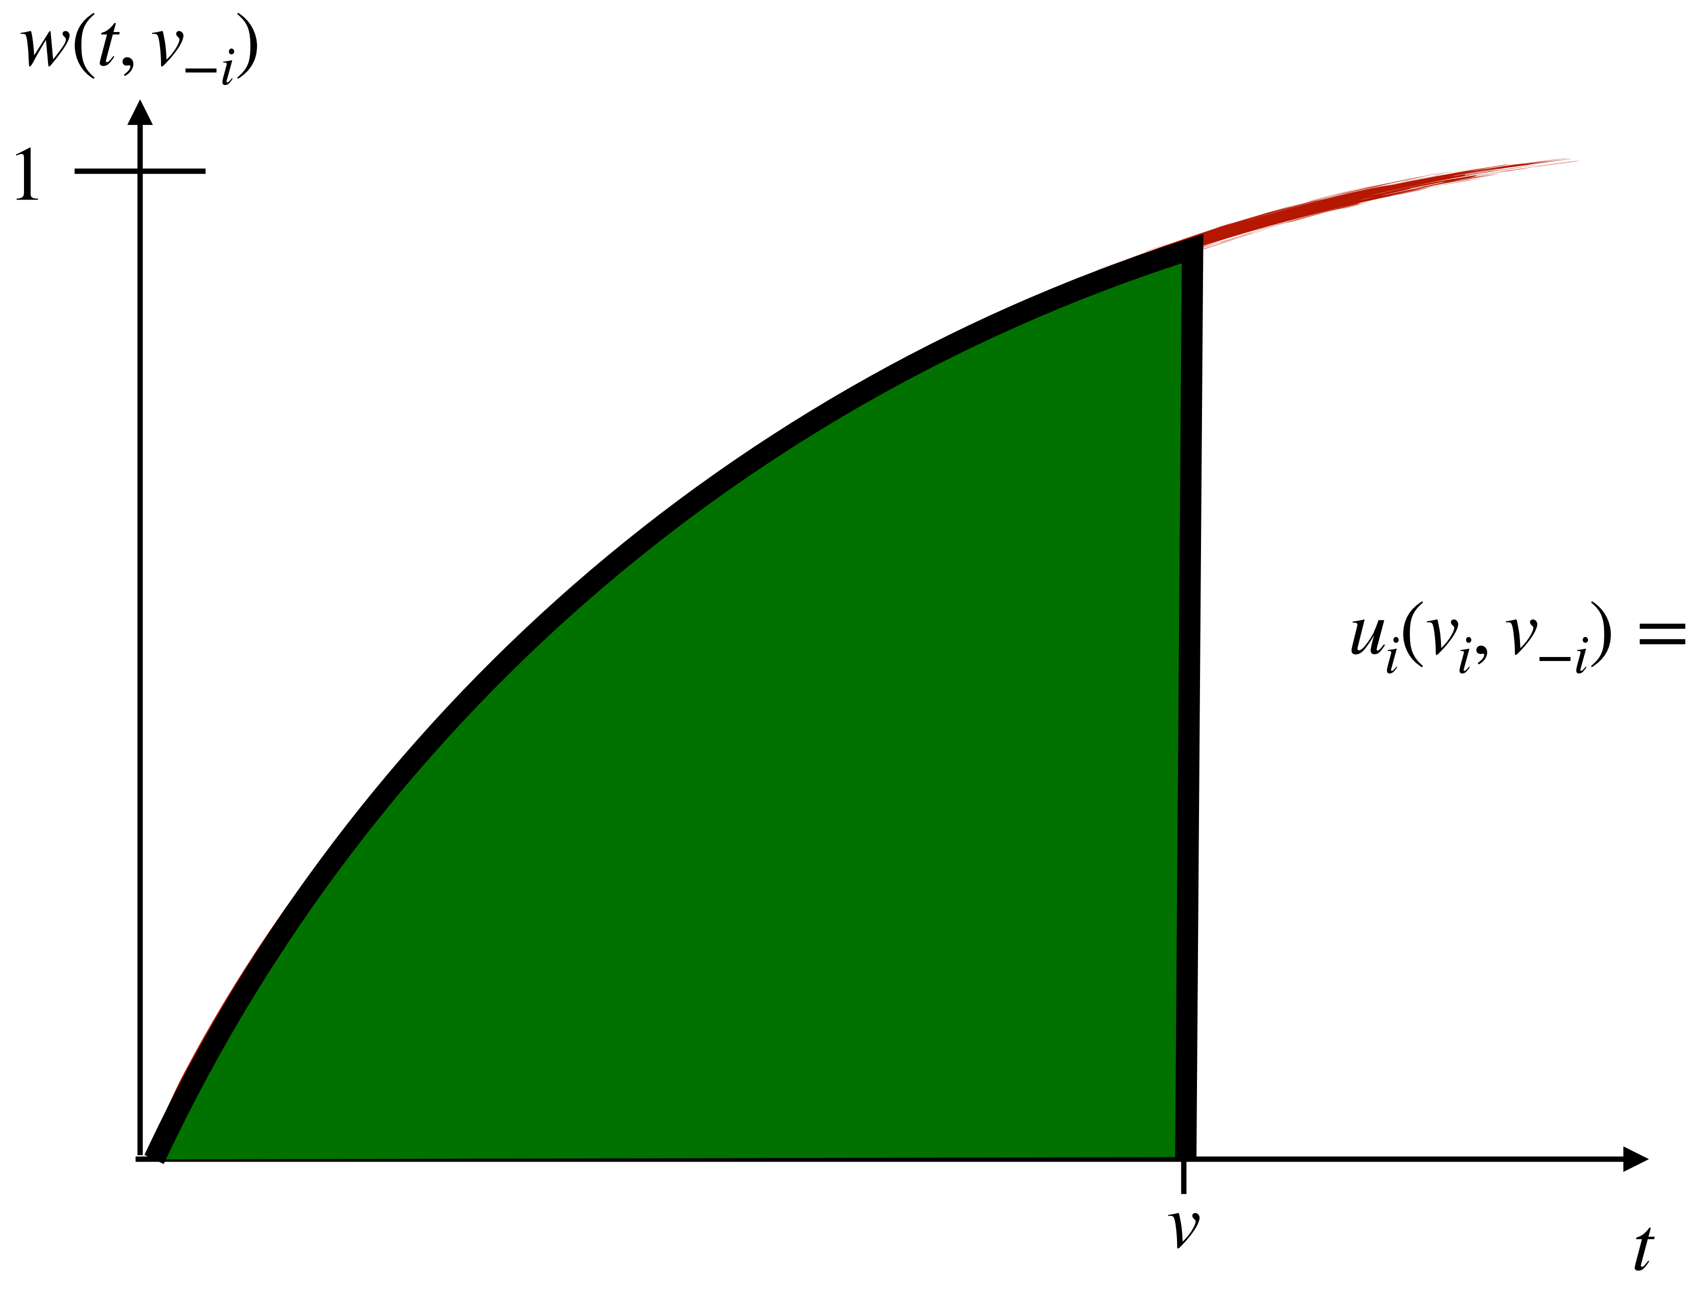

$$u_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

Some pictures

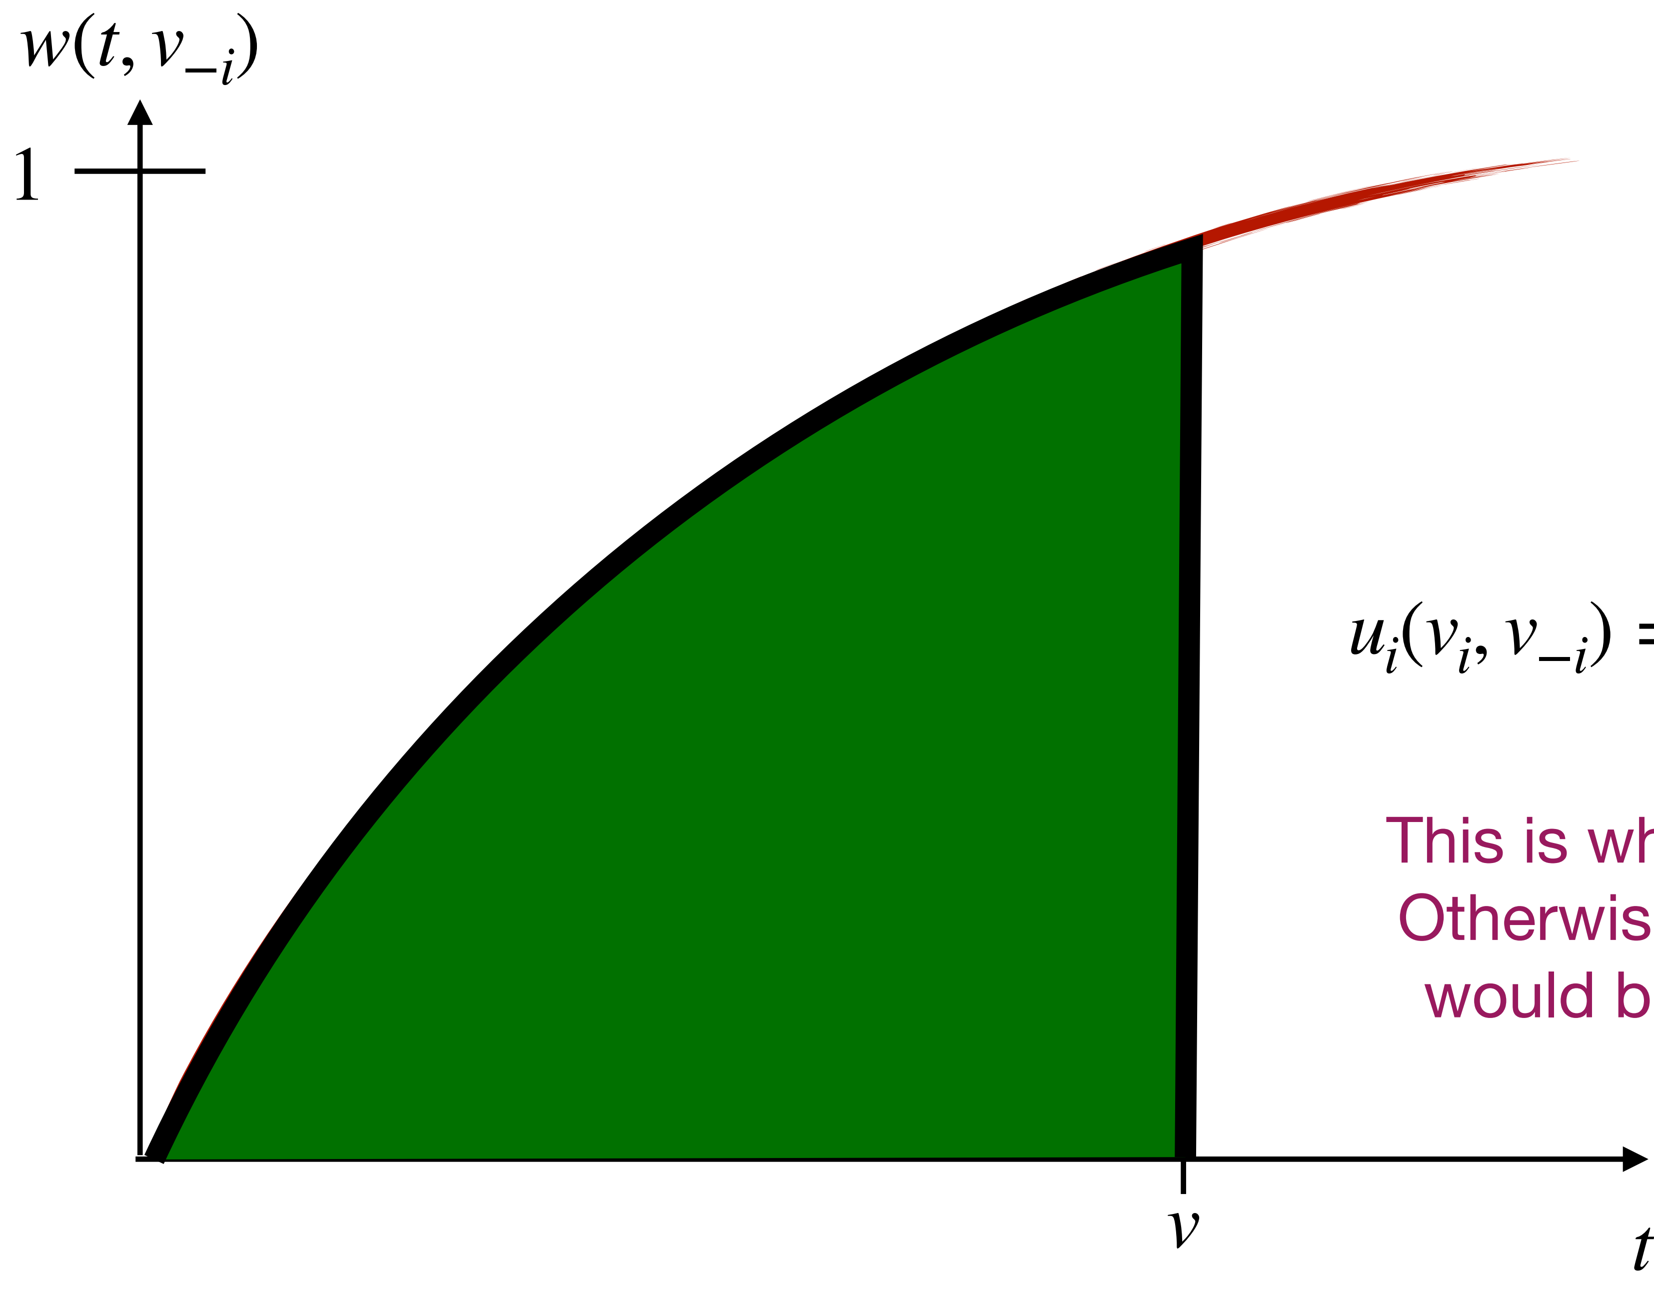

$$u_i(v_i, v_{-i}) = v_i \cdot w_i(v_i, v_{-i}) - p_i(v_i, v_{-i})$$

This is why monotonicity is important.
Otherwise the area under the integral
would be larger when misreporting!

More general mechanisms?

More general mechanisms?

<u>All the mechanisms that we have seen so far are of the following form:</u>

More general mechanisms?

All the mechanisms that we have seen so far are of the following form:

- The agents declare their values/preferences up front,

More general mechanisms?

<u>All the mechanisms that we have seen so far are of the following form:</u>

- The agents declare their values/preferences up front,

- The mechanism chooses an outcome (and payments) based on this declarations.

More general mechanisms?

<u>All the mechanisms that we have seen so far are of the following form:</u>

- The agents declare their values/preferences up front,

- The mechanism chooses an outcome (and payments) based on this declarations.

These mechanisms are called *direct revelation* mechanisms.

More general mechanisms?

All the mechanisms that we have seen so far are of the following form:

- The agents declare their values/preferences up front,

- The mechanism chooses an outcome (and payments) based on this declarations.

These mechanisms are called *direct revelation* mechanisms.

We could have more complicated mechanisms which e.g., interact with the agents in rounds, ask them questions, present them with tasks etc?

More general mechanisms?

All the mechanisms that we have seen so far are of the following form:

- The agents declare their values/preferences up front,
- The mechanism chooses an outcome (and payments) based on this declarations.

These mechanisms are called *direct revelation* mechanisms.

We could have more complicated mechanisms which e.g., interact with the agents in rounds, ask them questions, present them with tasks etc?

Crucially, could these mechanisms achieve things that truthful direct revelation mechanisms cannot?

What are truthful mechanisms really?

What are truthful mechanisms really?

A direct revelation mechanism has the following two properties:

What are truthful mechanisms really?

A direct revelation mechanism has the following two properties:

(1) For every valuation/preference profile, the corresponding game has a dominant strategy equilibrium (DSE).

What are truthful mechanisms really?

A direct revelation mechanism has the following two properties:

(1) For every valuation/preference profile, the corresponding game has a dominant strategy equilibrium (DSE).

(2) At this DSE, every agent truthfully reports her true value.

More general mechanisms

More general mechanisms

We would still like our mechanisms to have DSE.

More general mechanisms

We would still like our mechanisms to have DSE.

But we don’t need the agents to be truth-telling at those DSE.

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

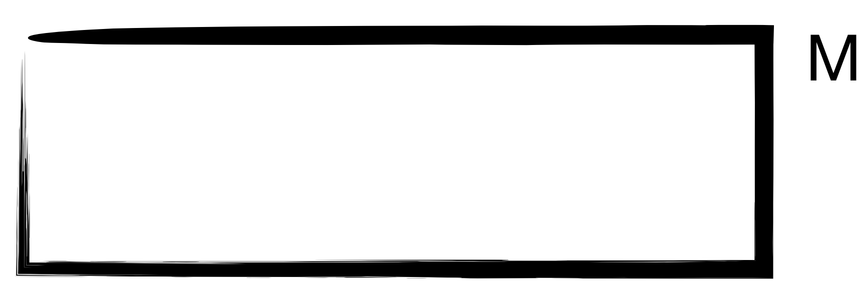

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

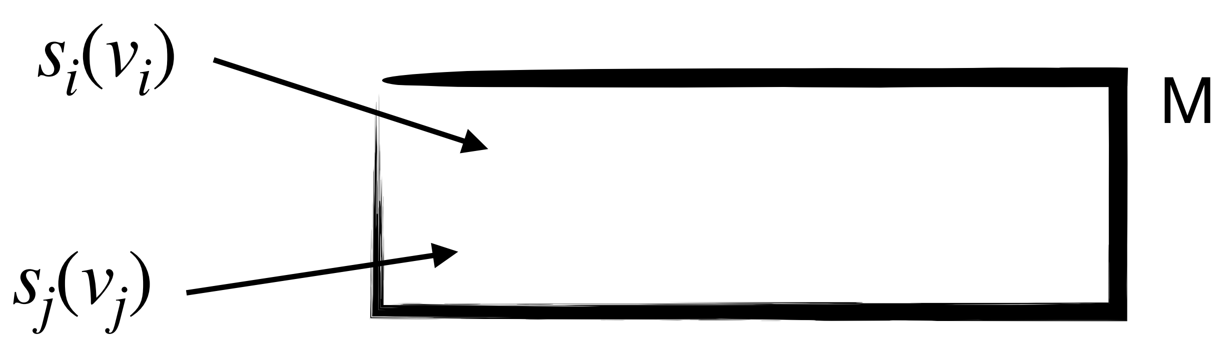

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

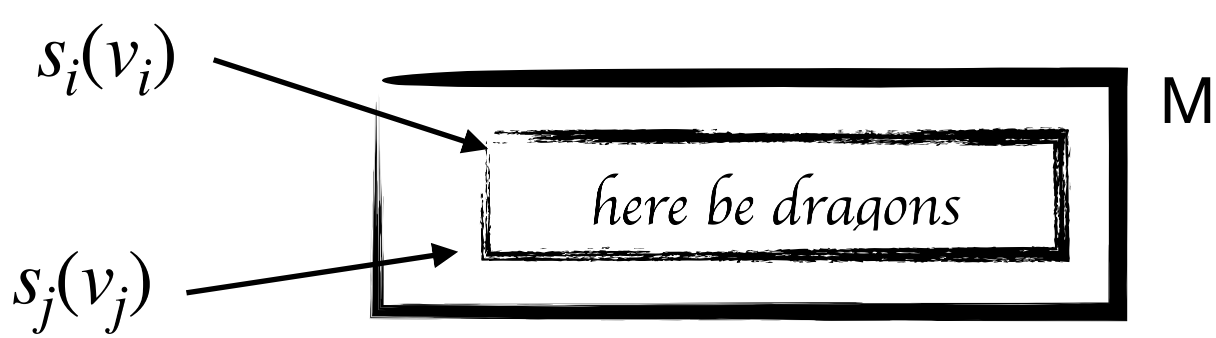

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

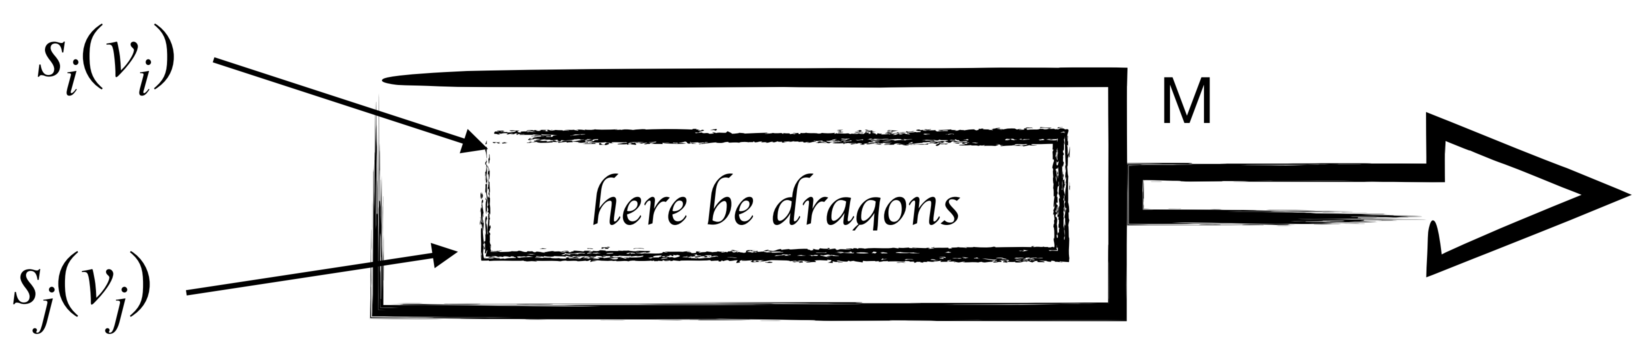

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

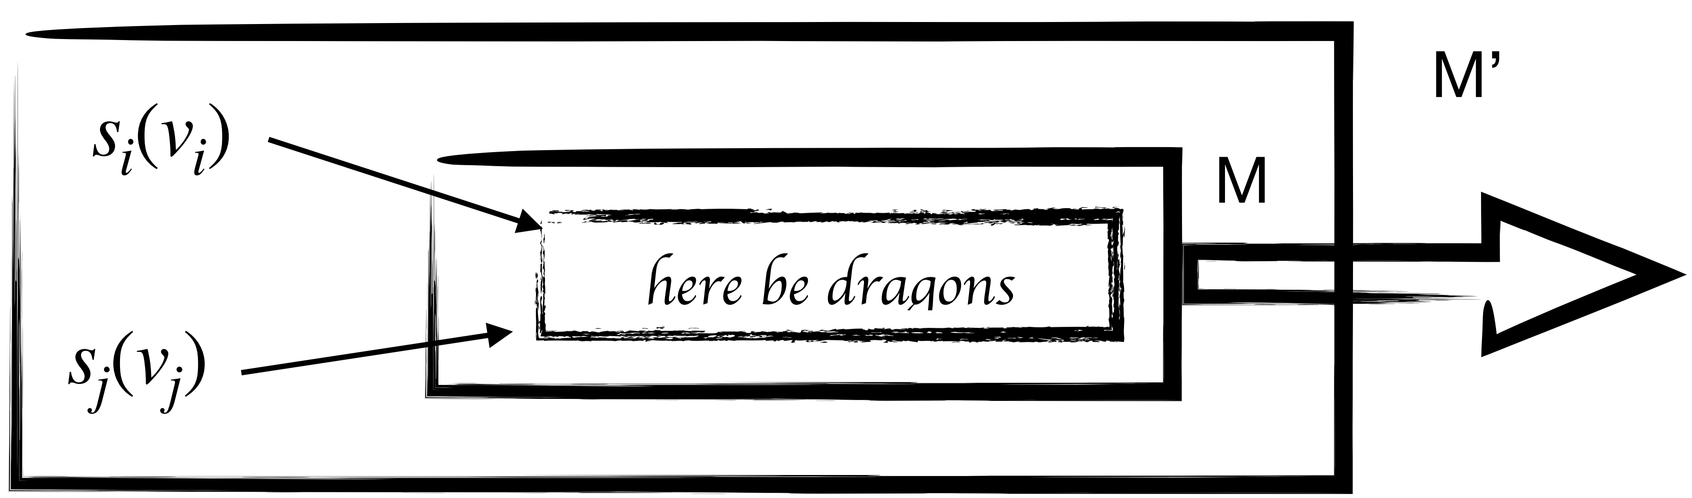

More general mechanisms

We would still like our mechanisms to have DSE.

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

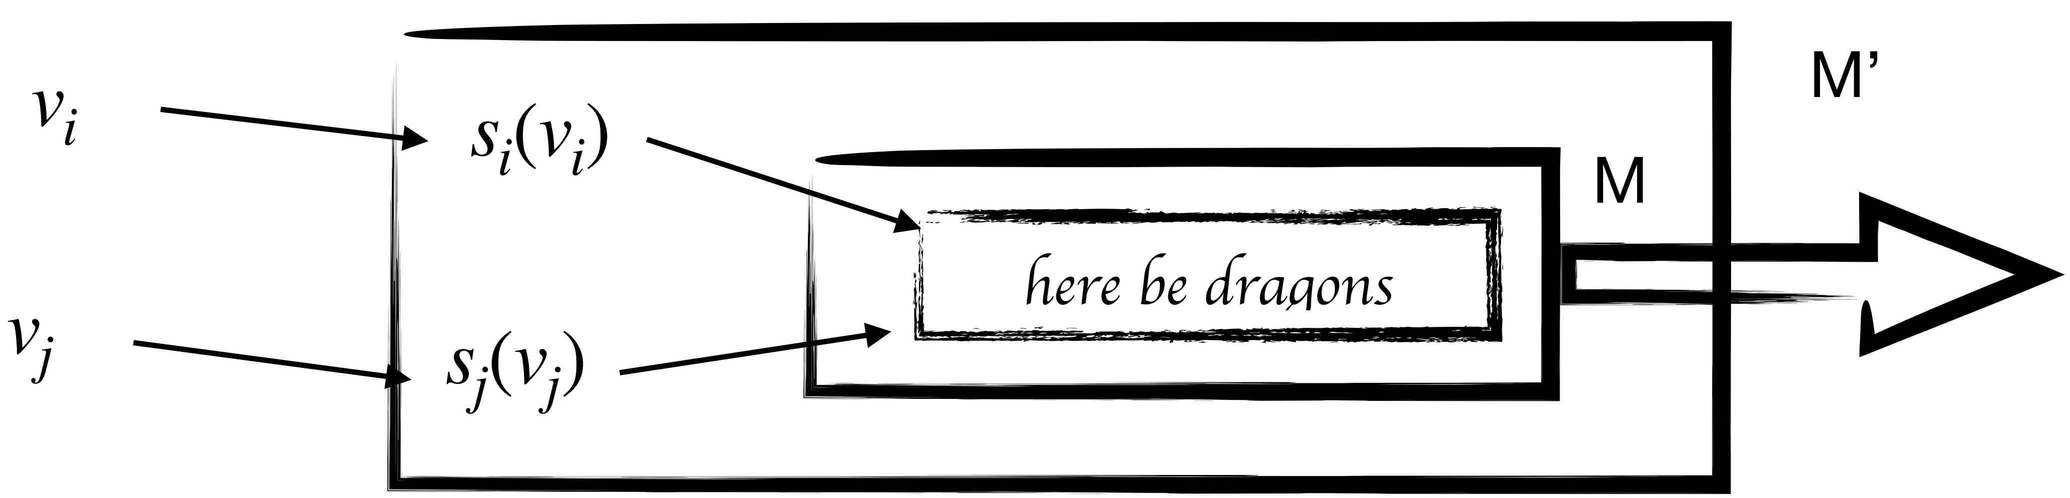

More general mechanisms

We would still like our mechanisms to have DSE. **What is we relax this requirement?**

But we don't need the agents to be truth-telling at those DSE.

And the mechanism does not have to be direct revelation.

Theorem (The Revelation Principle): If there is an *arbitrary* mechanism in which there is a DSE, then there is a *direct revelation truthful mechanism* (which uses the same payments and those obtained in the DSE).

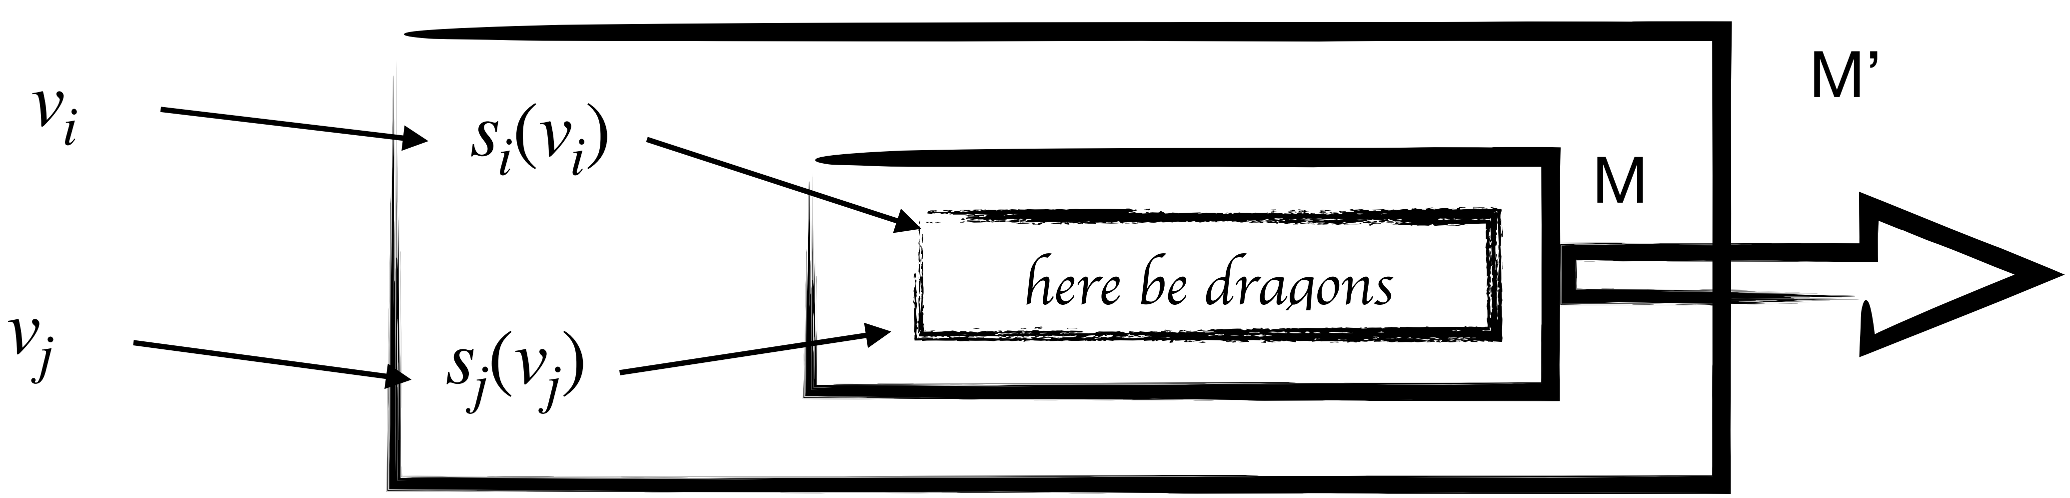

Beyond Truthful Mechanisms

Beyond Truthful Mechanisms

Maybe we are asking too much by requiring our mechanisms to be truthful.

Beyond Truthful Mechanisms

Maybe we are asking too much by requiring our mechanisms to be truthful.

Maybe we should "let people play" and use game theory to evaluate what will happen.

Beyond Truthful Mechanisms

Maybe we are asking too much by requiring our mechanisms to be truthful.

Maybe we should "let people play" and use game theory to evaluate what will happen.

If we insist on DSE, truthfulness is wlog (in settings where the Revelation Principle holds).

Beyond Truthful Mechanisms

Maybe we are asking too much by requiring our mechanisms to be truthful.

Maybe we should "let people play" and use game theory to evaluate what will happen.

If we insist on DSE, truthfulness is wlog (in settings where the Revelation Principle holds).

But we might be ok with just a mixed Nash equilibrium or a pure Nash equilibrium if it exists.

Beyond Truthful Mechanisms

Maybe we are asking too much by requiring our mechanisms to be truthful.

Maybe we should "let people play" and use game theory to evaluate what will happen.

If we insist on DSE, truthfulness is wlog (in settings where the Revelation Principle holds).

But we might be ok with just a mixed Nash equilibrium or a pure Nash equilibrium if it exists.

Or an appropriate notion for games with uncertainty (stay tuned).

Beyond Truthful Mechanisms

Maybe we are asking too much by requiring our mechanisms to be truthful.

Maybe we should "let people play" and use game theory to evaluate what will happen.

If we insist on DSE, truthfulness is wlog (in settings where the Revelation Principle holds).

But we might be ok with just a mixed Nash equilibrium or a pure Nash equilibrium if it exists.

Or an appropriate notion for games with uncertainty (stay tuned).

Could mechanisms with good Nash equilibria outperform truthful mechanisms?

Beyond Truthful Mechanisms

Maybe we are asking too much by requiring our mechanisms to be truthful.

Maybe we should "let people play" and use game theory to evaluate what will happen.

If we insist on DSE, truthfulness is wlog (in settings where the Revelation Principle holds).

But we might be ok with just a mixed Nash equilibrium or a pure Nash equilibrium if it exists.

Or an appropriate notion for games with uncertainty (stay tuned).

Could mechanisms with good Nash equilibria outperform truthful mechanisms?

Sometimes…