Frost and Frenzy:

Weather Risk, Forward Contracts and Tulipmania*

Peter Koudijs

25 March 2026

Abstract:

This paper proposes a fundamentals-based explanation for Tulipmania when tulip prices increased nearly eightfold between November 1636 and early February 1637. Tulip bulbs were planted in fall, only lifted from the ground in spring, and traded as forward contracts in between. Bulbs sustain damage under sub-freezing soil temperatures and the payment and delivery of forward contracts were conditional on bulbs being in good health, with sellers bearing the risk of damage. December 1636 and January 1637 were relatively cold, at the 13th percentile of historical winters, which historically predicted a cold February. Evidence from field experiments suggests that such temperatures would have caused significant damage to bulbs, reducing the expected future supply of healthy bulbs and driving up contract prices. In the presence of short-selling constraints, the drop in expected supply likely had an amplified effect on prices through belief heterogeneity. In line with this mechanism, price increases were largest among lower-end tulip varieties, where disagreement about aesthetic value was arguably greatest. When temperatures rose unexpectedly in early February, the anticipated supply shock failed to materialize and prices collapsed.

Keywords: Financial bubbles, Supply restrictions, Heterogeneous beliefs, Financial history

JEL codes: G13, G41, N23

* Peter Koudijs: NYU Stern and Erasmus School of Economics (pk3079@stern.nyu.edu). The author declares that he has no relevant or material financial interests related to the research described in this paper. I thank seminar and conference participants at … for feedback. I am grateful to Alexi Savov and Bruce Tuckman for valuable comments and suggestions. I thank Jiachen Jin, David Lever and Dorine Smit-Maarschalkerweerd for excellent research assistance. I am very grateful to Jan de Vries for generously sharing the data on canal closures used in the paper, and to Anne Goldgar for generously sharing her notes which provide additional prices for tulip bulbs. Any errors are my own.

In popular imagination, Tulipmania (1636–1637) has become nearly synonymous with irrational asset bubbles driven by herd behavior. How, indeed, could rational individuals pay the equivalent of four or more annual salaries of a skilled craftsman for a single tulip bulb?[1] The literature has sought to rationalize these high valuation *levels* by appealing to the exceptional cultural and aesthetic value of rare tulip varieties, high wealth inequality, and the fact that bulbs produce offsets capable of maturing into full bulbs, thereby generating future profits (Krelage 1942; Garber 2000; Goldgar 2007). What remains puzzling, however, is the price *dynamics*: on average, tulip bulb prices increased nearly eightfold between November 1636 and early February 1637 before collapsing shortly thereafter. None of the proposed rational explanations can account for so rapid a rise and so sudden a reversal (e.g. Van der Veen 2009).

This paper argues that these puzzling price dynamics reflect a fundamental factor that has so far been overlooked: weather conditions. A defining feature of the tulip trade is that bulbs remained in the ground for most of the year. They were planted in the fall and were only lifted from the ground after flowering in May or June. Consequently, all winter transactions took the form of forward contracts for future delivery. While in the ground, bulbs were exposed to damage risk, including from sub-freezing soil temperatures. Such damage manifests as reduced shoot weight and a lower probability of flowering and producing offsets, thereby diminishing bulb value. Delivery was therefore contingent on the bulb being in good condition, verified at the time of flowering. This contractual structure implies that the seller bore the risk of pre-delivery damage, while the buyer effectively held a state-contingent claim, with delivery and payment conditional on the bulb surviving in a verifiably good condition.

As long as demand is not perfectly elastic, this contingent-claim structure predicts higher contract prices in response to cold weather. Colder winters raise the probability that bulbs sustain damage, reducing the expected future supply of healthy bulbs and increasing the value of contracts written on surviving bulbs. Although damaged bulbs may recover over time, they do not contribute to the contemporaneous supply of tradeable bulbs because their quality is difficult to verify – particularly

[1] The popular association runs deep. In the film Wall Street: Money Never Sleeps, Gordon Gekko declares that "Tulipmania is the greatest bubble story of all time," claiming that a single bulb once fetched the price of a canal house in Amsterdam. Knowledge@Wharton similarly describes it as "a frenzy driven almost entirely by speculation, where prices detached from any real underlying value" (Benartzi and Puntoni 2025). The comparison in the text is based on the sale of a Generael der Generaelen van Gouda – an established variety – for f. 1,500 at the height of the boom; wage data are from De Vries and Van der Woude (1995, p. 707).

when they fail to produce a healthy flower. This mechanism implies a tight link between winter weather conditions and tulip prices.

Because asset prices are forward-looking, tulip prices should respond not only to realized weather conditions but also to expectations about future weather. Notably, a cold December and January historically predicted a higher likelihood of continued cold in February. Consequently, winter tulip prices should not only reflect contemporaneous damage to bulbs but also anticipated future damage, with continued price movements as adverse weather conditions persist.

The risk of temperature-induced damage to tulip bulbs would have been readily apparent to participants in the tulip trade. Economic life in the 17th century, situated in the middle of the little ice age, was broadly sensitive to weather conditions (De Vries and Van der Woude 1995, p. 37-40). Professional growers in particular, who constituted a large share of the market (Krelage 1942; Goldgar 2007), would have known from direct experience how consequential weather could be for the health of tulip bulbs.

Using historical weather data based on canal closures due to frost, the paper finds that December 1636 and January 1637 were relatively cold, falling at the 13th percentile of all winters between 1632 and 1757. The largest price increases of tulip bulb contract prices occurred during a period of prolonged frost in January 1637. Historically, these conditions predicted continued cold in February at the 28th percentile. Using experimental field evidence from Van der Valk (1970), this would have predicted a weight loss of tulip bulbs of approximately 12.5%. Because most tulip bulbs were planted in sandy soils with low groundwater levels – conditions that substantially amplify the detrimental effects of low temperatures (AVvB 1930; Krelage 1942, pp. 96–97, 103) – realized damage was likely greater than this baseline estimate.

In reality, the cold weather did not persist. The frost broke at the end of January, and February proved warmer than anticipated, registering at the 52nd percentile of historical winters, well above the 28th percentile expected during the height of the cold. Using Van der Valk (1970), this corresponds to a realized weight loss of only 7.5%, significantly lower than the expected 12.5%. Consistent with this mechanism, tulip prices began to decline after January 31, coinciding with the thawing temperatures, falling by as much as 25% within five days before trade collapsed entirely (Roman 1637a, p. 17). That realized damage was indeed limited is confirmed by a contemporaneous inspection record: on March 20, 1637, a buyer examining a large number of

tulips found only some shoots to be "not in very good health," warranting a price reduction of only 2.6% of the contract value (Posthumus 1927, p. 54).

It is possible that changes in expected supply amplified price dynamics through heterogenous beliefs. When market participants disagree about an asset's fundamental value and short selling is restricted, prices reflect the valuation of the marginal investor, who may be relatively optimistic. A contraction in asset supply leads to higher prices because the asset is held by fewer investors, pushing the marginal trader toward a more optimistic type (Miller 1977; Harrison and Kreps 1978; Scheinkman and Xiong 2003; Hong, Scheinkman, and Xiong 2006). Because tulips were luxury goods whose value was closely tied to fashion, market participants can plausibly be expected to have held heterogeneous beliefs about the future popularity of tulips in general and of specific varieties in particular. Furthermore, because tulip bulbs were largely non-fungible, naked short selling was impossible: sellers could only take forward positions on specific bulbs planted in the ground (Goldgar 2007, p. 218). A weather-induced decline in expected supply could therefore have amplified price dynamics through this heterogeneous-beliefs channel.

To test this hypothesis, the paper examines whether the price dynamics differed systematically across tulip varieties likely to be subject to greater or lesser belief disagreement. Evidence from the art economics literature suggests that valuation dispersion – and hence belief disagreement – is lower at the higher end of the market than at the lower end (e.g. Chambers, Dimson, and Spaenjers 2020), a pattern that plausibly extends to other goods with high intrinsic aesthetic value including tulip bulbs.[2] We therefore compare price dynamics between high-end collector varieties, for which valuations were likely more widely agreed upon, and more common or newly introduced low-end varieties, for which beliefs about future popularity were likely more dispersed. Consistent with the heterogeneous-beliefs mechanism, the price increase and subsequent reversal were almost four times as large for lower-end varieties than for high-end ones.

This paper is related to a large literature on the formation of financial bubbles.[3] Its key contribution lies in advancing a plausible explanation for one of the most famous historical bubbles. In

[2] Goldgar (2007) explicitly draws parallels between investing in art and tulip bulbs during the 17th century, with both conveying a high aesthetic and cultural value.

[3] See Scherbina (2013) for an overview; Blanchard and Watson (1982) and Tirole (1985) on the existence of "rational bubbles;" Allen and Gale (2000) on agency problems in financial intermediation; and Shleifer and Vishny (1997), Mitchell, Pulvino and Stafford (2002), Abreu and Brunnermeier (2003), Brunnermeier and Nagel (2004) and Temin and Voth (2004) on limits to arbitrage and rational agents "riding a bubble."

particular, it is the first paper to provide evidence that changes in fundamentals – specifically, weather conditions and changes in expected supply – contributed to the rise and decline of forward prices of tulip bulbs during Tulipmania. More broadly, the paper contributes to an empirical literature connecting asset float, short-sale constraints, heterogeneous beliefs, and financial bubbles (e.g. Ofek and Richardson 2003, Lamont and Thaler 2003, Mei, Scheinkman and Xiong 2009, and Xiong and Yu 2011). Relative to this literature, the paper exploits exogenous variation in weather conditions to identify the supply-price link, providing cleaner identification than prior work relying on lockup expirations or market segmentation. Further, whereas the existing literature mostly documents supply expansions that deflate bubbles, this paper links changes in expected supply to both the rise and the collapse of prices, tracing the full arc of the bubble to a single underlying shock and its unexpected reversal. Finally, it provides a direct cross-sectional test of the heterogeneous-beliefs mechanism, showing that price swings were nearly four times as large for tulip varieties that plausibly exhibited the most belief disagreement about future popularity.

The paper builds on a long literature studying Tulipmania. Posthumus (1926, 1927, 1929, 1934) reproduces an extensive overview of primary sources, while Krelage (1942) provides a first synthesis of the episode. Both argue that the entry of new non-professional investors, in combination with the commodification of certain varieties, fueled speculation, with prices becoming increasingly detached from fundamentals until the pool of new buyers was exhausted. Dash (1999) similarly emphasizes speculative dynamics, arguing that the institutional setting of the trade, organized through the so-called "colleges," facilitated this process. Van der Veen (2009) emphasizes the tightly knit and geographically concentrated social network, arguing that Tulipmania arose as a consequence of an information cascade or herding. Goldgar (2007), the most thorough and complete historical study to date, emphasizes the exceptional cultural and aesthetic value attributed to certain tulip varieties, which rationalizes high valuation levels, especially in an environment of high wealth inequality. However, even in this detailed analysis, the sharp rise and subsequent collapse in prices remains ultimately unexplained, with Goldgar relying on a rational bubble interpretation to account for it.[4] Goldgar also notes the widespread refusal to honor contracts after the price collapse and the limited degree of enforcement. It is unclear whether this

[4] "There was nothing intrinsically crazy (…) about buying a product it was clear one could sell on at a higher price. The unsustainability of the price was not predictable, and if a crash happened, it was not necessarily to be foreseen. (…) Tulipmania was only irrational after the fact; if the market had held, it would have been supremely sensible to invest one's money in tulip bulbs." Goldgar (2007, p. 227).

was anticipated ex ante – forward contracts were longstanding financial instruments, firmly embedded in an established legal framework (Van Tielhof 2002, Gelderblom and Jonker 2004, 2005; Petram 2014). Further, the refusal to honor tulip forward contracts came at a substantial reputational cost. This makes a pure risk-shifting explanation unlikely (Goldgar 2007, pp. 228–229, 283–284).

Other contributions argue that Tulipmania can be explained by fundamentals alone, but each relies on demonstrably incorrect assumptions. Garber (1990) argues that tulip prices gradually declined due to natural supply growth through offshoots and that there was no bubble. This disregards a transaction for multiple bulbs from May 1, 1637 that indicates a dramatic price collapse, which is supported by the historical record (Roman 1637c, p. 12; Goldgar 2007, p. 378). Garber's argument also cannot account for the rapid run-up in December 1636 and January 1637. Thomson (2006) suggests that tulip forward prices were strike prices of call options, finding no bubble in spot prices. Beyond failing to explain why strike prices would rise sharply while spot prices would not, his argument rests on a single transaction recorded at $f70$, which he interprets as a spot sale, which was actually recorded at $f260$, the lower figure being a typographical error in Posthumus's (1929) English translation.[5] There is also no evidence that this transaction's contractual form differed from others in the same record that Thomson classified as call options. McClure and Thomas (2017) argue that traders underestimated tulip supply and revised their expectations upward upon observing shoots emerging in early February 1637. This mechanism cannot explain the price run-up, and the timing is implausible: given soil temperatures in Haarlem well below 4°C through January 31, shoots are unlikely to have emerged at that point.[6]

The remainder of the paper proceeds as follows. Section I provides historical background. Section II develops the paper's argument into testable hypotheses. Section III describes the sources and data construction for the tulip bulb price indices and weather data. Section IV presents the empirical evidence, linking weather conditions to tulip forward prices and documenting larger price increases for lower-end varieties. Section V concludes.

[5] The original Dutch source states: "Op den 9en December 1636 ghekocht van N.N. een Gheel ende Root van Leyden...voor de somma van tweehondert ende tsestich gulden [two hundred and sixty]" (Roman 1637b, p. 14).
[6] Sustained soil temperatures above 4°C are required for bulbs to sprout (Cheng et al. 2022). Benschop (1980) shows that even in average Dutch winters, shoots only reach sufficient height to emerge above ground in mid-March.

I. Historical background[7]

A. The Dutch Tulip trade in 1636-1637

Tulip bulbs were prized in early 17th century Holland for the beauty of their flowers. Tulip collecting occupied the same cultural world as art collecting and was seen as a way to enhance one's social status. Collectors applied an identical critical vocabulary to flowers and paintings alike, and transactions sometimes blurred the boundary entirely: bulbs were exchanged directly for paintings by old masters.

As Goldgar (2007) observes, the beauty of a tulip was not self-evident but culturally constructed. Starting in the late 16th century, collectors began to develop a hierarchy of varieties. The most coveted varieties displayed vivid stripes and flames of contrasting color – a phenomenon caused by a virus that would not be identified until the twentieth century and that could not be reliably reproduced – which heightened their perceived rarity and value. The exact hierarchy of such different varieties was a matter of fashion. According to Goldgar (2007, p. 122), "flowers went in and out of mode in this period," with newer varieties capturing the attention of collectors and finding their way into this hierarchy. It would take until the 18th century before tulip standards were solidified into an unambiguous set of "florists' flowers."

During the winter of 1636–1637, market participants were not trading physical bulbs for immediate delivery, but rather claims on bulbs that remained planted in the ground. Most transactions concerned future delivery after flowering, with payment likewise deferred. Bulbs were traded as specific named varieties, often priced by weight, which was indicative of future reproductive potential through offsets. The contractual structure of tulip transactions closely resembled forward contracts. Buyers and sellers typically agreed on a price in winter, with settlement contingent on later delivery of a healthy bulb. The actual contracts written shared the language of other forward contracts that were routinely written on assets such as Dutch East India stock or commodities (Van Tielhof 2002, Gelderblom and Jonker 2004, 2005; Petram 2014), and made the buyer fully liable for execution of the contract (Roman 1637b, p. 15). However, because neither goods nor money usually changed hands at the time of contracting, performance still relied on trust, reputation, and honor.

[7] Unless otherwise noted, this section is based on Goldgar (2007).

The forward structure of the market allowed a single bulb to change hands multiple times over the course of the winter, while it remained in the ground where it was originally planted. Yet this did not make shorting easy. Each transaction was tied to a specific, physically located bulb, making naked short-selling – selling a forward on a bulb one did not own – effectively impossible (Goldgar, pp. 182–183, 194, 203–206). An active forward market thus coexisted with substantial shorting constraints.

Tulip trading was organized around inns and taverns that functioned as recurring market venues, hosting "colleges" where participants met regularly to negotiate and conclude transactions. Sales proceeded through semi-private bargaining mediated by knowledgeable arbiters, or through auction-like mechanisms with ritualized bidding practices. Price discovery thus occurred within socially embedded networks rather than anonymous competitive markets, with trusted intermediaries playing a central role in valuing heterogeneous bulbs and establishing fair prices.

B. The boom and bust in forward Tulip bulb prices

Tulip prices rose over the late fall and early winter of 1636-1637 but accelerated sharply in late January. While some elite varieties had commanded high prices for decades, the final weeks before the collapse saw rapid appreciation across a wider range of bulbs, including lower-end varieties. Contemporary observers emphasized the suddenness of this acceleration, which marked a distinct shift from earlier, more gradual price increases (e.g. Roman 1637b, p. 24). Section IV.A has more precise quantitative estimates of the development of tulip bulb prices.

According to the historical record, trading broke down in Haarlem on February 3, 1637, when members of a college were unable to sell tulip bulbs at previously prevailing peak prices. News of this collapse spread rapidly through the interconnected colleges, causing demand to evaporate almost immediately (Roman 1637a, p. 17). Although some transactions continued briefly such as in Alkmaar, confidence in the broader market collapsed within days with a “large and extreme fall in prices” (Posthumus 1927, pp. 7, 49; Goldgar 2007, pp. 230-231). In response, local representatives met in Amsterdam on February 23, 1637 to discuss what measures could be taken to avoid “the damages, yes utmost ruin” that some would expect from the “extreme rise in prices” and subsequent crash (Posthumus 1927, p. 49). No firm agreement was reached (Goldgar 2007, p. 236)

When forward contracts came due in May or June of 1637, many buyers refused to take delivery. Parties initially sought to resolve disputes privately. Provincial authorities declined to void contracts, and a States of Holland proclamation of April 27 kept all transactions in full force, directing local magistrates to encourage friendly settlements rather than impose them. In the end, many contracts were settled for a fraction of the contracted price.

Whether this outcome was anticipated ex ante is far from obvious. Like any other forward contract traded in this period, tulip contracts imposed full liability on the buyer (Roman 1637b, p. 15). But even setting aside legal enforcement, refusing delivery carried severe reputational consequences in a commercial society built on mutual credit and personal honor. The tulip trade operated within small, tightly interconnected networks bound by family ties, shared religion, and overlapping business dealings. A failure to honor a contract was immediately visible to the entire community and threatened one's creditworthiness well beyond the flower trade. As Goldgar notes, there was little that could be worse for a merchant's reputation than refusing to pay his debts, and the Haarlem authorities specifically warned that defaults would bring ruin and shame upon persons of good name (Goldgar 2007, pp. 228–229, 237). Taken together, these considerations suggest that the limited ex post enforcement of tulip contracts was not a primary driver of the bubble.

C. Forward contracts and weather risk

Importantly, final payment and delivery of a forward contract on a tulip bulb were contingent on it producing a healthy flower.[8] If a bulb sustained damage, whether through frost, disease, or any other form of deterioration, it was less likely to flower in future seasons and to produce offsets, directly diminishing its intrinsic value (e.g. Posthumus 1927, pp. 54, 61-62). Beyond this intrinsic loss, a damaged bulb that failed to flower also could not be delivered under the terms of the contract. Delivery required the buyer to be present to inspect the flower and confirm it matched what had been contracted: a seller could not lift bulbs without the buyer witnessing the flower first, since otherwise the buyer could dispute the identity or quality of what was delivered. This verification requirement made flowering itself a practical precondition for enforceable delivery, a norm confirmed in 1639 when three Haarlem tulip experts testified that "the custom is, and always has been, that when someone sells a tulip bulb when growing, whether as a Gouda, or otherwise,

[8] For concrete examples of this contingent contract structure and further discussion, see Posthumus 1927, pp. 13, 21, 33, 39, 54, 81, 85; Posthumus 1934, p. 231; Krelage 1942, pp. 61–66; Goldgar 2007, pp. 62–63, 198, 286, 289–290.

and the tulip bulb does not bring forth such a flower as the seller has sold to the buyer, that the buyer is not obliged to receive it" (Posthumus 1927, p. 85; Goldgar 2007, p. 286; Van der Veen 2009, p. 22). Importantly, even though damaged bulbs could recover over time, they did not add to the contemporaneous supply of a given variety precisely because, in the absence of a healthy flower, their identity could not be verified. Because damage resulted in fewer offsets, it also reduced the growth of future supply.

An important source of damage risk was the weather and subfreezing soil temperatures in particular. This would have been readily apparent to participants in the tulip trade. Tulipmania was situated in the middle of the little ice age, when temperatures in the Dutch Republic were uncommonly low (Shabalova and van Engelen 2003; Collet et al. 2025). Economic life, which in large part centered on agriculture, was broadly sensitive to weather conditions (e.g. De Vries and Van der Woude 1995, p. 37-40). Professional growers in particular, who constituted a large share of the market (Krelage 1942; Goldgar 2007), would have known from direct experience how consequential cold weather could be for the health of tulip bulbs.

The link between subfreezing soil temperatures and tulip bulb damage is well-established in horticultural literature studying tulip bulbs in the Netherlands. AVvB (1930, p. 6) documents that the severe winter of 1928–1929 caused substantial bulb damage. Damage was especially severe for bulbs planted in sandy soils with low groundwater levels – conditions that closely describe the tulip-growing region around Haarlem during Tulipmania (Krelage, pp. 96–97, 103; Goldgar 2007, pp. 137, 158-159, 165; Kaartenatlas Haarlem, pp. 15, 19). KAVvB (1963) and CvdTH (1985) record similar damage during the harsh winters of 1962–1963 and 1984–1985. Van der Valk and Geers (1969, 1970) and Van der Valk (1970) provide experimental evidence on the effects of subzero soil temperatures on bulb damage, establishing the critical role of both the intensity and total duration of freezing.

II. Hypotheses

The particular contract structure through which forwards on tulip bulbs were traded, together with the sensitivity of bulb health to winter weather conditions, and the inherent uncertainty about what tulip varieties might be in vogue next season, imply the following four testable hypotheses.

Hypothesis 1: Tulip contract prices are positively associated with current and expected cold winter temperatures.

Colder conditions increase the probability of bulb damage and reduce expected future supply. If demand for tulip bulbs is not fully elastic, this should predict higher contract prices on surviving bulbs. This effect is amplified to the extent that cold temperatures are serially correlated: if cold weather today signals colder weather tomorrow, prices should reflect not only contemporaneous damage but also the anticipated continuation of adverse conditions.

Hypothesis 2: Tulip contract prices fall following an unexpected positive weather shock.

Because forward prices are set on the basis of expected future supply, a sudden and unanticipated rise in winter temperatures – which limits damage to planted bulbs and restores expected supply – should produce an immediate decline in contract prices.

Hypothesis 3: An unexpected positive weather shock materially reduces realized bulb damage relative to what had been anticipated.

Damage to tulip bulbs is proxied by the weight of shoots growing from a bulb. Using agronomic estimates of the relationship between soil temperatures and shoot weight, it is possible to calculate the difference in damage arising from expected and realized winter temperatures.

Hypothesis 4: Price increases and subsequent reversals are larger for tulip varieties subject to greater valuation dispersion.

The first three hypotheses examine supply shocks and forward prices through a standard demand-supply framework. Hypothesis 4 identifies an additional channel: heterogeneous beliefs about fundamental value.

The ultimate value of any tulip variety was determined by future fashion – an inherently subjective quantity about which market participants were likely to disagree. When short selling is restricted, such disagreement can push prices above the market consensus, as pessimists are unable to counteract the buying pressure of optimists (Miller 1977). A contraction in supply amplifies this effect. The resulting price increase should be largest for varieties subject to the greatest valuation

dispersion – precisely those for which disagreement about future fashion was most pronounced (Harrison and Kreps 1978; Scheinkman and Xiong 2003; Hong, Scheinkman, and Xiong 2006).

III. Data

A. Tulip prices

To track the development of Tulipmania, the paper uses all reliable prices for tulip bulb contracts that can be found in the historical archives and contemporary sources. Data are drawn from Roman (1637a-c), Posthumus (1927, 1934), and Goldgar (2007), supplemented by the original research notes for Goldgar (2007), generously provided by Anne Goldgar. The information from Posthumus and Goldgar come primarily from notarial acts, with Goldgar providing additional prices from contemporary bills of sale. Roman (1637a-c) provides three pamphlets published contemporaneously with the mania. Where Roman's prices can be cross-checked against the notarial record in Posthumus and Goldgar, they appear reliable (Van der Veen 2009).

To summarize price developments, the paper constructs a price index. The primary challenge is that individual varieties, often with different weights, were traded infrequently and at different points in time. A price index can still be estimated through a hedonic regression of log contract prices on period fixed effects, controlling for bulb-level characteristics and variety fixed effects. Specifically, we estimate:

$$\log(P_{i,t}) = \sum_{t \neq 15Nov} \beta_t \tau_t + X'_{i,t}\Gamma + \psi_i + \varepsilon_{i,t} \quad (1)$$

where $P_{i,t}$ is the price of a forward contract on a bulb (bulbs) of variety i traded in period t, τ_t are period fixed effects, $X'_{i,t}$ are control variables, specifically the log(weight) of bulbs before planting and dummies for the type of bulb traded (offset, single bulb, pound of bulbs, or multiple pounds of bulbs), and ψ_i are variety fixed effects. Only varieties with at least two transactions are included. The estimated price index is the exponential of the period fixed effects β_t, normalized to 1 in November 1636, the omitted category. To correctly infer the statistical uncertainty of these estimates, standard errors are clustered at the variety level.

Periods t are defined as intervals wide enough to ensure a sufficient number of transactions for estimation: $t \in$ {November 1-30, December 1-31 (1636), January 1-15, 16-27, 28, 30, 31, February 1, 2, 3, 4, 5, 5-9, May 1 (1637)}. The granularity increases towards the peak of the market as more

transactions were being reported in the sources. Because transactions are sometimes recorded for a multi-day span rather than a single date, we create 10,000 datasets in which such transactions are allocated randomly across the relevant days using a uniform distribution. Reported coefficients are means across these datasets, and 95% confidence intervals are constructed from the mean standard deviation across datasets.

B. Weather

Direct temperature records from the 17th century are scarce and of uncertain reliability. The paper therefore exploits episodes in which the canal between Amsterdam and Haarlem – where most traded tulip bulbs were planted – froze over. The canal was served by a horse-drawn passenger barge, the dominant form of inter-city travel in the Dutch Republic, which carried passengers, light goods, and mail on scheduled departures several times daily. Operating under municipal concession, its overseers were required to maintain detailed revenue accounts recording each trip. When the canal froze, barges could not run and no fares were collected, leaving gaps in the record. Figure A.1 in the Online Appendix has a map of the canal and Figure A.2 shows the barge service as depicted by contemporaneous painter Salomon van Ruysdael.

Although a crude proxy for general weather conditions, a frozen canal is directly indicative of the sub-freezing soil temperatures that would have damaged planted tulip bulbs. The period also falls within the little ice age, when average temperatures were considerably lower than today and waterways froze with sufficient regularity to make canal closures a meaningful signal – as vividly attested by the frequent depiction of ice skating in contemporary Dutch paintings. Figure A.2 in the Online Appendix shows an example by Adam van Breen, depicting a frozen canal along which barges would ordinarily have run.

Data on canal closures were generously provided by Jan de Vries (de Vries 1977, 1978, 1980). Monthly data are available for 1632–1757, while daily data are available only from 1656 onwards. To construct a daily series for the winter of 1636–1637, the paper draws on daily closure records from the canal between Arnhem and Nijmegen, located approximately 90 km to the east, which are available for 1610–1639. The Amsterdam–Haarlem canal is assumed closed on any day the Arnhem–Nijmegen canal reported no traffic. Where a given month contains additional Amsterdam–Haarlem closure days not accounted for by this method, they are allocated to days on which the Arnhem–Nijmegen canal reported little traffic, immediately following a no-traffic day.

To map canal closures to actual temperatures, we use direct thermometric observations for 1706–1756 from the Dutch Royal Meteorological Institute (KNMI), which provides records for Delft-Rijnsburg (1706–1734), Haarlem (1735–1742), and Zwanenburg (1738–1760). Zwanenburg, situated near the midpoint of the Amsterdam–Haarlem canal, serves as the baseline series, which is extended backwards using readings from the other stations via OLS regressions estimated over overlapping periods to adjust for differences in baseline temperatures across stations.

The weather stations above record air temperature at approximately two meters above ground. To convert these into soil temperatures, we draw on data from AVvB (1930), which reports simultaneous air and soil temperature observations collected at the government research laboratory in Lisse – approximately 15 km south of Haarlem – over the winter of 1928–1929. Soil temperature is measured at a depth typical for tulip bulbs, under a cover of common reed. Because there is no historical evidence that growers in the 1630s used such insulating cover, our soil temperature estimates are best interpreted as an upper bound on the warmth experienced by bulbs in the ground. The relation between soil (S_t) and air temperatures (A_t) is estimated as

$$S_t = \alpha + \beta A_t + \varepsilon_t, \tag{2}$$

The coefficient β captures how strongly the current day's air temperature drives soil temperature Regression estimates are reported in Table A.1 in the Online Appendix. The R^2 of the regression is 0.62. Figure A.3 shows the in-sample fit. The coefficient estimates $\{\alpha, \beta\}$ are then used to map historical air temperatures into soil temperatures.

Finally, we employ a simulation exercise to estimate a distribution of plausible soil temperatures for each day of the winter of 1636–1637. Each canal closure episode in 1636–1637 is matched to episodes of similar length from 1706–1756, when daily temperatures are observed directly, using brackets of ±1 day for episodes up to 10 days, ±2 days for episodes of 11–20 days, and so on. From these matched episodes, we generate 10,000 possible temperature series by randomly drawing and combining closure episodes, weighting each candidate by the similarity of its winter to that of 1636–1637. Similarity is measured by the fraction of days the canal remained open over the full winter; weights are defined as the inverse of the squared difference in this fraction from that of 1636–1637 and are normalized to sum to one for each episode. For each matched episode, we take the soil temperatures of the closure period together with the three preceding days, reflecting the fact that ice typically requires several consecutive days of sub-freezing temperatures to form.

Figure A.4 in the Online Appendix illustrates our approach, showing the distribution of plausible soil temperatures implied by the canal closures during the winter of 1636-1637.

IV. Empirical analysis

A. Tulip bulb prices and the weather

Hypothesis 1 predicts that tulip bulb contract prices should be positively associated with current and expected cold winter temperatures.

Figure 1 presents the estimated price index from November 1636 through May 1637. The index reveals a dramatic run-up concentrated in the second half of January. Prices were relatively stable throughout November and December 1636, rising modestly before accelerating sharply after mid-January 1637. By late January and early February, prices had reached approximately eight times their November level. Confidence intervals are wide at the peak, reflecting substantial cross-sectional dispersion in the prices of individual varieties, but remain well above the November baseline throughout.

This price increase unfolded against a backdrop of sustained cold. Figure 2 compares canal closures over December and January 1636–1637 to all winters from 1632 to 1757, revealing considerable year-to-year variation in winter severity. The vertical line marks the observation for December 1636 through January 1637: the canal was frozen for 60% of days, placing that winter at the 13th percentile of the historical distribution. While not the most extreme winter on record, the cold was substantial.

Crucially, a cold December and January was historically predictive of a cold February. Figure 3 presents a binscatter of the fraction of days the canal was open in December and January against the corresponding fraction in February. The relationship is positive and precisely estimated, with a t-statistic of 5.90: winters that were cold in the first two months of the season tended to remain cold into February. Participants in the tulip trade would therefore have had good reason to expect, at the height of the January cold, that adverse conditions would persist – and to price contracts accordingly.

B. Weather conditions and shocks to expected supply

Hypothesis 2 predicts that tulip contract prices should fall following an unexpected positive weather shock.

Figure 4 plots the distribution of the fraction of days in February that the Amsterdam–Haarlem canal was open, across all winters from 1632 to 1757. The two vertical lines indicate the expected and realized February fractions for 1637. The expected line is based on a linear regression of the February fraction on the December–January fraction, evaluated at the observed 1636–1637 December–January value; it corresponds to approximately the 28th percentile of the historical distribution. Around it, the figure reports the estimated normal distribution taking into account both sampling variation and parameter uncertainty. The realized line, representing actual February 1637 conditions, falls at approximately the 52nd percentile, well above the expected value. The conditional probability of observing a February as mild as the one in 1637 – given the severity of the cold in the preceding December and January – was only 16.4%, confirming that the warming came as a substantial surprise.

Figure 1 shows that prices fell dramatically and abruptly after January 31, declining by as much as 25% within five days before trade collapsed entirely after February 5. By May 1, 1637 – the first day on which transactions were recorded after the market collapse – prices had fallen well below their peak and even below their November 1636 level.

Figure 5 overlays the daily estimated canal closures with the estimated price index. The shaded areas indicate the days the canal was closed and the solid line gives the cumulative closure days. The price and cumulative closure day series move in close correspondence, suggesting a causal relation. The dramatic price increase materialized during a prolonged period of canal closures, while the price decline started precisely on the day the canal reopened, marking the end of the cold spell and the arrival of a warmer-than-expected February.

C. Weather conditions and expected tulip bulb damage

Hypothesis 3 predicts that the unexpected positive weather shock materially reduced realized bulb damage relative to what had been anticipated.

Damage to tulip bulbs is typically proxied by the weight loss of the shoots they produce in spring. Van der Valk (1970, Figure 2) shows that this is strongly predictive of the probability of flowering. As discussed in Section I.B, a bulb that fails to produce a healthy flower cannot have its variety verified, rendering any forward contract written on it void. Further, such bulbs therefore do not contribute to the supply of healthy, deliverable tulips. Also, weight loss predicts fewer future offshoots, also limiting supply further into the future.

To compare realized against expected tulip damage, the paper constructs a series of counterfactual canal closure days under the assumption that the cold winter continued into February 1637. Figure 6 shows that, across all winters between 1632 and 1757, canal closures were significantly more likely in the first half of February. Accordingly, the expected canal closure days for February 1637 from Figure 4 are allocated to the beginning of the month. The distribution of expected daily soil temperatures is then calculated using the same methodology described in Section III.B used for realized soil temperatures.

To translate soil temperatures into estimates of bulb damage, we draw on Van der Valk (1970), who reports results from controlled experiments conducted near Lisse – approximately 15 km south of Haarlem – documenting the relationship between the duration of exposure to subfreezing soil temperatures and the weight of the emerging shoot, where a healthy shoot weighs 20 grams. Damage is measured as the weight shortfall relative to this benchmark. Figure A.5 in the Online Appendix presents his data and shows that the relationship between the length of exposure to a given subfreezing soil temperature and bulb damage is approximately linear. This implies that expected tulip damage can be calculated as the simple product of damage per day and the number of days. Van der Valk (1970) provides estimates for prolonged exposure to soil temperatures of −1°C and −3°C. Damage estimates for soil temperatures between 0°C and −1°C, or between −1°C and −3°C, are obtained through interpolation.

Figure 7 presents the daily distributions of plausible tulip bulb damage during the winter of 1636–1637 under both expected and realized temperature paths. To arrive at these distributions, expected tulip damage is calculated using the soil temperatures in each of the 10,000 simulated temperature paths using either realized or expected canal closure days.

The divergence between expected and realized damage in the beginning in February is substantial. Under the expected temperature path – reflecting the 28th percentile February prediction given the observed December–January cold – the median expected weight of planted bulbs is about 17.5 grams, implying a weight loss of 12.5% by the end of winter. This is likely an underestimate of the expected damage: the soil temperature estimates are based on data where the ground was covered with an insulating layer of common reed, yet there is no historical evidence that 17th-century growers used this practice, suggesting actual soil temperatures were lower than calculated. Moreover, in 1636–1637 most bulbs were planted in gardens south of Haarlem characterized by

sandy soil and low groundwater levels; conditions that AVvB (1930, p. 6) identifies as producing more severe temperature-induced damage, based on observations from the exceptionally cold winter of 1928–1929. Even at the estimated 12.5%, the expected shortfall in the supply of healthy, deliverable bulbs would warrant a significant increase in forward contract prices for surviving ones.

Under the realized temperature path – reflecting the milder February that actually occurred at the 52nd percentile – the median realized weight of planted bulbs is approximately 18.5 grams, implying a loss of 7.5%, well below the expected 12.5%. Anecdotal evidence is consistent with limited realized damage. On March 20, 1637 a buyer inspected the shoots of a large number of tulips priced at f. 11,700. Some shoots were found to be "not in very good health," but the price was only reduced by f. 300 (-2.6%), suggesting that overall damage was limited (Posthumus 1927, p. 54). The improvement in expected supply conditions in early February 1637 thus provides a direct fundamental rationale for the observed price collapse.

D. Heterogeneity across varieties

Hypothesis 4 predicts that price increases and subsequent reversals should be larger for tulip varieties subject to greater valuation dispersion. When market participants disagree about fundamental value, short selling is constrained, and asset supply is restricted, prices reflect the valuation of the marginal investor – who will be relatively optimistic. A contraction in supply pushes prices higher still: as the asset is held by fewer investors, the marginal trader shifts toward a more optimistic type.[9]

The tulip market of 1636–1637 fits this setting well. Shorting was effectively restricted by the non-fungibility of bulbs: forward contracts could only be written on specific physical bulbs (Goldgar 2007, p. 218). And as discussed in Section I.A, tulip valuations were culturally determined, with no established consensus about which varieties were most desirable – making heterogeneous beliefs about future popularity both natural and likely.

To assess the relative importance of this channel, we exploit variation in the degree of belief disagreement likely to prevail across tulip varieties. Following the logic of the heterogeneous-beliefs framework, a supply contraction should produce larger price increases for varieties for

[9] Back (2010, pp. 357-365) provides a clear textbook treatment of the basic theoretical result.

which beliefs about fundamental value are more dispersed and the few optimists setting the price are relatively more optimistic than for a variety for which beliefs are less dispersed. Evidence from the art economics literature suggests that experts' valuations are more dispersed at the low end of the art market than at the high end. Chambers, Dimson, and Spaenjers (2020, p. 507, Figure 3) show that expert appraisals show more dispersion for lower-end paintings. Similarly, Aubry et al. (2023, p. 818, Figure 9, Panel C) shows that the relation between auction house presale estimates and final hammer prices is noisier for the lower end of the market. This pattern plausibly extends to tulip varieties, given that – as discussed in Section I.A – market participants attributed to tulip bulbs the same aesthetic value as to other art forms, paintings in particular.

To divide tulips into low and high-end varieties, the paper uses the variety fixed effect estimates ψ_i of Eqn. (1) that provide the price level of a traded variety corrected for the general movement in prices during the winter of 1636-1637. Figure 8 shows the distribution of these fixed effects and documents that there is a distinct set of high-end tulip varieties that have a log-valuation that is clearly outside the rest of the distribution. The tulip varieties of this elite group all correspond to prestigious varieties that were well-established before November 1636: the Admirael Liefkens, Van der Eyck, and Van Enckhuysen, and the Bruyne Purper, Generael der Generaelen van Gouda, Scipio, and Vice Roy (Krelage 1942, Goldgar 2007) Belief disagreement about their intrinsic worth was likely smaller than for the newer, less well-established varieties at the lower end of the market.

Figure 9 presents separate price indices for high-end and lower-end varieties estimated from:

$$\log(P_{i,t}) = \sum_{t \neq 15Nov} \beta_t \tau_t + \sum_{t \neq 15Nov} \gamma_t \tau_t \times L_i + X'_{i,t}\Gamma + \psi_i + \varepsilon_{i,t}, \quad (3)$$

which follows the same notation as Eqn. (1) and where L_i is a dummy for lower end varieties. The estimation is the same as for Eqn. (1). Coefficients β_t capture the price index for high-end varieties, and $\beta_t + \gamma_t$ the index for low-end varieties. The figure plots the 95% confidence intervals for the estimated coefficients γ_t centered on $\beta_t + \gamma_t$. As such, they indicate the statistical significance of the difference between high and lower end varieties.

Results are consistent with a shock to expected supply amplifying price dynamics through heterogenous beliefs. While both types of varieties experienced price increases during the winter

of 1636–1637, the average magnitude of these movements was almost four times as large for lower-end varieties. The difference in price dynamics between the two groups is highly statistically significant, especially around the peak.

V. Conclusion

This paper offers a fundamentals-based explanation of one of history's most famous financial bubbles, Tulipmania, when forward prices on tulip bulbs increased almost eightfold during the winter of 1636-1637. The paper's explanation is rooted in the interaction between weather risk and contract characteristics. Bulbs remained in the ground all winter and were traded as forward contracts, with delivery contingent on the bulb producing a healthy flower – the only means by which a variety could be verified. Because sellers bore the risk of pre-delivery damage, anything reducing the expected supply of healthy, verifiable bulbs would raise the value of surviving contracts. Frost was precisely such a threat: subfreezing soil temperatures reduce bulbs' shoot weight and significantly lower the probability of flowering.

The winter of 1636–1637 began under icy conditions. Considering the number of days the Amsterdam–Haarlem canal was frozen over, December and January were relatively cold, falling at the 13th percentile of historical winters. Historically, a cold December and January predicted a cold February, giving market participants good reason to anticipate continued frost, which would seriously impair the health of tulip bulbs. Using experimental agronomic data and simulated soil temperatures derived from canal closures, the paper estimates that expected shoot weight loss reached approximately 12.5%, implying a substantial contraction in the future supply of verifiable, deliverable bulbs.

February, however, proved far milder than the preceding cold had led participants to expect – at the 52nd percentile of historical winters against an expected 28th percentile, an outcome with only a 16.4% probability. The Amsterdam–Haarlem canal reopened on February 1, and prices began falling the same day, dropping by as much as 25% within five days before the market collapsed entirely. This is consistent with implied shoot weight loss falling to around 7.5%, well below the anticipated 12.5%, meaning the expected contraction in the supply of verifiable, deliverable bulbs was far less severe than feared. A contemporaneous inspection of a large group of tulips corroborates this conclusion, finding damage warranting a price reduction of just 2.6%.

The effect of the expected supply shock on prices was likely amplified by heterogeneous beliefs. Because tulip forward contracts could only be written on specific existing bulbs, short-selling was restricted. In this environment, a contraction in expected supply pushes the marginal buyer toward a more optimistic type, amplifying price increases. Consistent with this, price increases were nearly four times as large for lower-end tulip varieties, where disagreement over aesthetic value and future popularity was plausibly greatest, than for established high-end collector varieties whose valuations likely commanded broader consensus.

Taken together, these findings suggest that the dramatic price increases and subsequent reversals of Tulipmania were not simply the product of irrational speculation. Price dynamics likely reflected rational responses to genuine uncertainty about an impending supply disruption – one that ultimately did not arrive. Behavioral factors may still have played a role, but not in the form of blind speculation: in the presence of shorting restrictions, the anticipated supply contraction likely interacted with fundamental disagreement about variety-level value to amplify the price dynamics.

References

Abreu, Dilip, and Markus K. Brunnermeier. 2003. "Bubbles and Crashes." Econometrica 71 (1): 173–204.

Algemeene Vereeniging voor Bloembollencultuur (AVvB). 1930. Onderzoek vorstschade: Rapport van de commissie, belast met het instellen van een onderzoek naar de schade, veroorzaakt door de weersomstandigheden gedurende den winter van 1928–1929.

Allen, Franklin, and Douglas Gale. 2000. "Bubbles and Crises." Economic Journal 110 (460): 236–255.

Aubry, Mathieu, Roman Kräussl, Gustavo Manso, and Christophe Spaenjers. 2023. "Biased Auctioneers." Journal of Finance 78 (2): 795–833.

Back, Kerry. 2010. Asset Pricing and Portfolio Choice Theory. Oxford: Oxford University Press.

Benartzi, Shlomo, and Stefano Puntoni. 2025. "Decoding the AI Bubble: Tulips, Dot-coms, or Something Else?" Knowledge@Wharton, December 16. https://knowledge.wharton.upenn.edu/article/decoding-the-ai-bubble-tulips-dot-coms-or-something-else/.

Benschop, M. 1980. "Growth and Development of Tulip CV 'Apeldoorn,' from Planting until Emergence." Acta Horticulturae 109: 189–196.

Blanchard, Olivier J., and Mark W. Watson. 1982. "Bubbles, Rational Expectations, and Financial Markets." In Crises in the Economic and Financial Structure, edited by Paul Wachtel, 295–315. Lexington, MA: Lexington Books.

Brunnermeier, Markus K., and Stefan Nagel. 2004. "Hedge Funds and the Technology Bubble." Journal of Finance 59 (5): 2013–2040.

Chambers, David, Elroy Dimson, and Christophe Spaenjers. 2020. "Art as an Asset: Evidence from Keynes the Collector." Review of Asset Pricing Studies 10 (3): 490–520.

Cheng, Chen, Liping Feng, J. F. Bienvenido Barcena, Weidong Yu, Gen Li, Zhenfa Li, and Caihua Ye. 2022. "A Growth Model Based on Standardized Growing Degree Days for Hydroponic Fresh Cut Tulip in Solar Greenhouses." European Journal of Horticultural Science 87 (4): 1–13.

Collet, Dominik, Sam White, Scott Bremer, Brita Brenna, Håkon Glørstad, Ingar Mørkestøl Gundersen, Heli Huhtamaa et al. 2025. "The Little Ice Age: The History and Future of a Traveling Concept." Wiley Interdisciplinary Reviews: Climate Change 16 (5): e70023.

Consulentschap van de Tuinbouw te Hoorn (CvdTH). 1985. Interne enquête vorstschade 1984–1985.

Dash, Mike. 1999. Tulipomania: The Story of the World's Most Coveted Flower and the Extraordinary Passions It Aroused. London: Gollancz.

de Vries, Jan. 1977. "Histoire du climat et économie: des faits nouveaux, une interprétation différente." Annales: Économies, Sociétés, Civilisations 32 (2): 198–226.

de Vries, Jan. 1978. Barges and Capitalism: Passenger Transportation in the Dutch Economy, 1632–1839. A.A.G. Bijdragen 21. Wageningen.

de Vries, Jan. 1980. "Measuring the Impact of Climate on History: The Search for Appropriate Methodologies." Journal of Interdisciplinary History 10 (4): 599–630.

de Vries, Jan, and Ad van der Woude. 1995. Nederland 1500–1815: De Eerste Ronde van Moderne Economische Groei. Amsterdam: Balans.

Garber, Peter M. 1990. "Famous First Bubbles." Journal of Economic Perspectives 4 (2): 35–54.

Garber, Peter M. 2000. Famous First Bubbles: The Fundamentals of Early Manias. Cambridge, MA: MIT Press.

Gelderblom, Oscar, and Joost Jonker. 2004. "Completing a Financial Revolution: The Finance of the Dutch East India Trade and the Rise of the Amsterdam Capital Market, 1595–1612." Journal of Economic History 64 (3): 641–672.

Gelderblom, Oscar, and Joost Jonker. 2005. "Amsterdam as the Cradle of Modern Futures Trading and Options Trading, 1550–1650." In The Origins of Value: The Financial Innovations that Created Modern Capital Markets, edited by William N. Goetzmann and K. Geert Rouwenhorst, 189–206. Oxford: Oxford University Press.

Goldgar, Anne. 2007. Tulipmania: Money, Honor, and Knowledge in the Dutch Golden Age. Chicago: University of Chicago Press.

Harrison, J. Michael, and David M. Kreps. 1978. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." Quarterly Journal of Economics 92 (2): 323–336.

Hong, Harrison, José A. Scheinkman, and Wei Xiong. 2006. "Asset Float and Speculative Bubbles." Journal of Finance 61 (3): 1073–1117.

KNMI (Royal Dutch Meteorological Institute). "Historic Observations of the Netherlands (Antieke Reeksen)." https://dataplatform.knmi.nl/dataset/antieke-reeksen-v1-0.

Koninklijke Algemeene Vereeniging voor Bloembollencultuur (KAVvB). 1963. Enquête naar het optreden van vorstschade 1963.

Krelage, Ernst Heinrich. 1942. Bloemenspeculatie in Nederland. Amsterdam: P. N. van Kampen.

Lamont, Owen A., and Richard H. Thaler. 2003. "Can the Market Add and Subtract? Mispricing in Tech Stock Carve-Outs." Journal of Political Economy 111 (2): 227–268.

McClure, James E., and David Thomas. 2017. "Tulip Mania: The Classic Story of the Dutch Tulip Speculation Scandal." Eastern Economic Journal 43 (2): 343–351.

Mei, Jianping, José A. Scheinkman, and Wei Xiong. 2009. "Speculative trading and stock prices: Evidence from Chinese AB share premia." Annals of Economics and Finance 10-2: 225-255.

Miller, Edward M. 1977. "Risk, Uncertainty, and Divergence of Opinion." Journal of Finance 32 (4): 1151–1168.

Mitchell, Mark, Todd Pulvino, and Erik Stafford. 2002. "Limited Arbitrage in Equity Markets." Journal of Finance 57 (2): 551–584.

Ofek, Eli, and Matthew Richardson. 2003. "DotCom Mania: The Rise and Fall of Internet Stock Prices." Journal of Finance 58 (3): 1113–1137.

Petram, Lodewijk. 2014. The World's First Stock Exchange. New York: Columbia University Press.

Posthumus, N. W. 1926. "De speculatie in tulpen in de jaren 1636 en 1637 (I)." Economisch-Historisch Jaarboek 12: 3–99.

Posthumus, N. W. 1927. "De speculatie in tulpen in de jaren 1636 en 1637 (II)." Economisch-Historisch Jaarboek 13: 3–85.

Posthumus, N. W. 1929. "The Tulip Mania in Holland in the Years 1636 and 1637." Journal of Economic and Business History 1 (3): 434–466.

Posthumus, N. W. 1934. "De speculatie in tulpen in de jaren 1636 en 1637 (III)." Economisch-Historisch Jaarboek 18: 229–240.

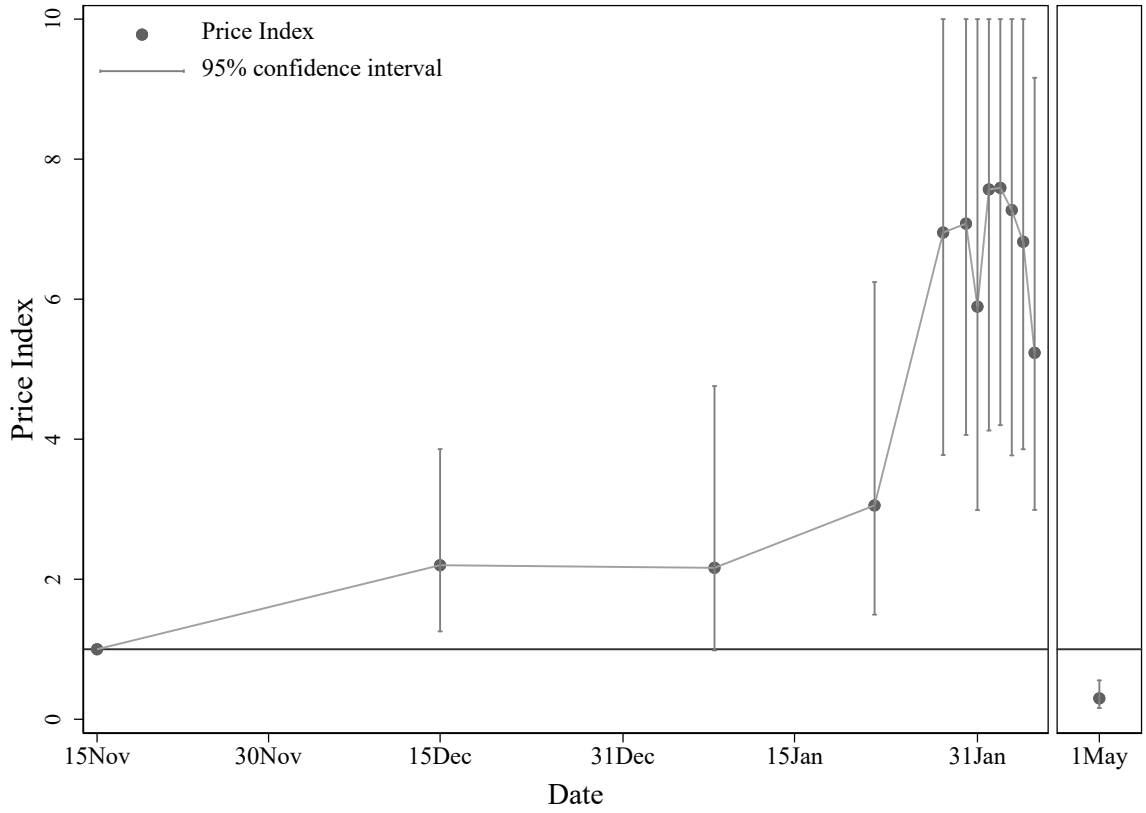
Roman, Adriaen. 1637a. Eerste Samenspraeck tusschen Waermondt en Gaergoedt, Nopende de opkomst ende ondergangh van Flora. Haarlem.

Roman, Adriaen. 1637b. Tweede Samenspraeck tusschen Waermondt en Gaergoedt, Nopende de opkomst ende ondergangh van Flora. Haarlem.

Roman, Adriaen. 1637c. Register van de Prijzen der Bloemen, Zijnde de Derde Samenspraeck tusschen Waermondt en Gaergoedt, Nopende de opkomst ende ondergangh van Flora. Haarlem.
Scheinkman, José A., and Wei Xiong. 2003. "Overconfidence and Speculative Bubbles." Journal of Political Economy 111 (6): 1183–1219.
Scherbina, Anna. 2013. Asset Price Bubbles: A Selective Survey. Washington, DC: International Monetary Fund.
Shabalova, Marina V., and Aryan FV Van Engelen. 2003. "Evaluation of a reconstruction of winter and summer temperatures in the Low Countries, AD 764–1998." Climatic Change 58 (1): 219-242.
Shleifer, Andrei, and Robert W. Vishny. 1997. "The Limits of Arbitrage." Journal of Finance 52 (1): 35–55.
Temin, Peter, and Hans-Joachim Voth. 2004. "Riding the South Sea Bubble." American Economic Review 94 (5): 1654–1668.
Thomson, Earl A. 2006. "Tulipmania: Fact or Artifact?" Public Choice 130 (1–2): 99–114.
Tirole, Jean. 1985. "Asset Bubbles and Overlapping Generations." Econometrica 53 (6): 1499–1528.
van der Valk, G. G. M. 1970. "Frost Injury to Flowerbulb Crops." Acta Horticulturae 23: 345–349.
van der Valk, G. G. M., and K. Geers. 1969. "Vorstschade in bloembolgewassen." Jaarverslag Laboratorium voor Bloembollenonderzoek Lisse, 111–112.
van der Valk, G. G. M., and K. Geers. 1970. "Vorstschade in bloembolgewassen." Jaarverslag Laboratorium voor Bloembollenonderzoek Lisse, 101.
van Tielhof, Milja. 2002. The 'Mother of All Trades': The Baltic Grain Trade in Amsterdam from the Late Sixteenth to the Early Nineteenth Century. Leiden: Brill.

Figures

Figure 1: Estimated price index, Nov. 1636 – May 1637.

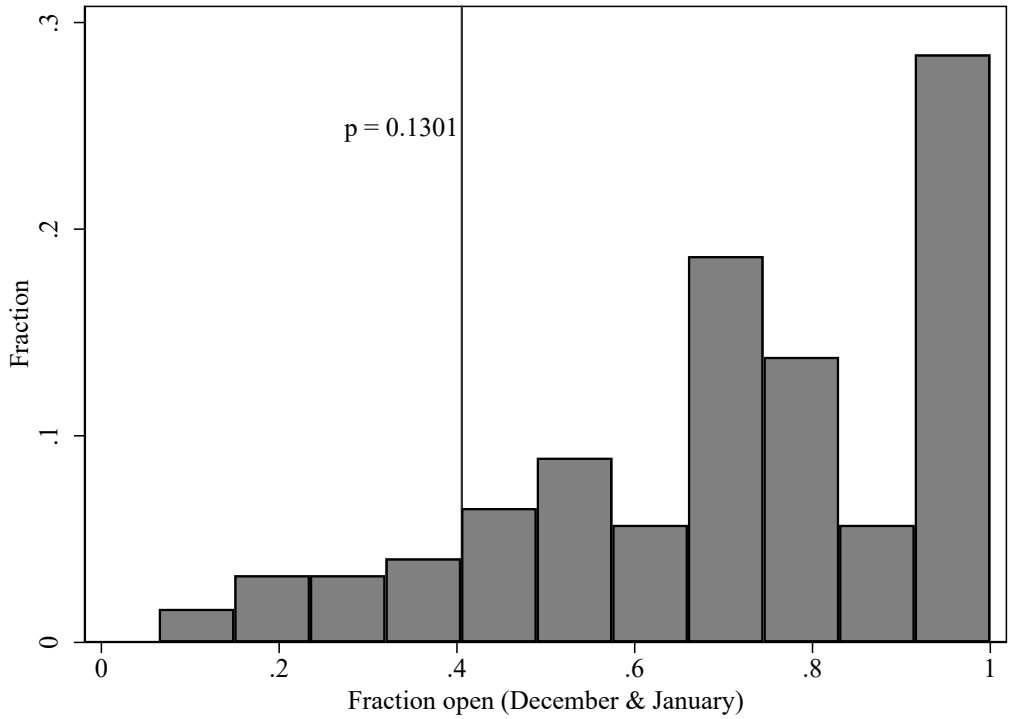


The graph reports the exponential of the time fixed effects β_t from:

$$\log(P_{i,t}) = \sum_{t \neq 15Nov} \beta_t \tau_t + X'_{i,t}\Gamma + \psi_i + \varepsilon_{i,t}$$

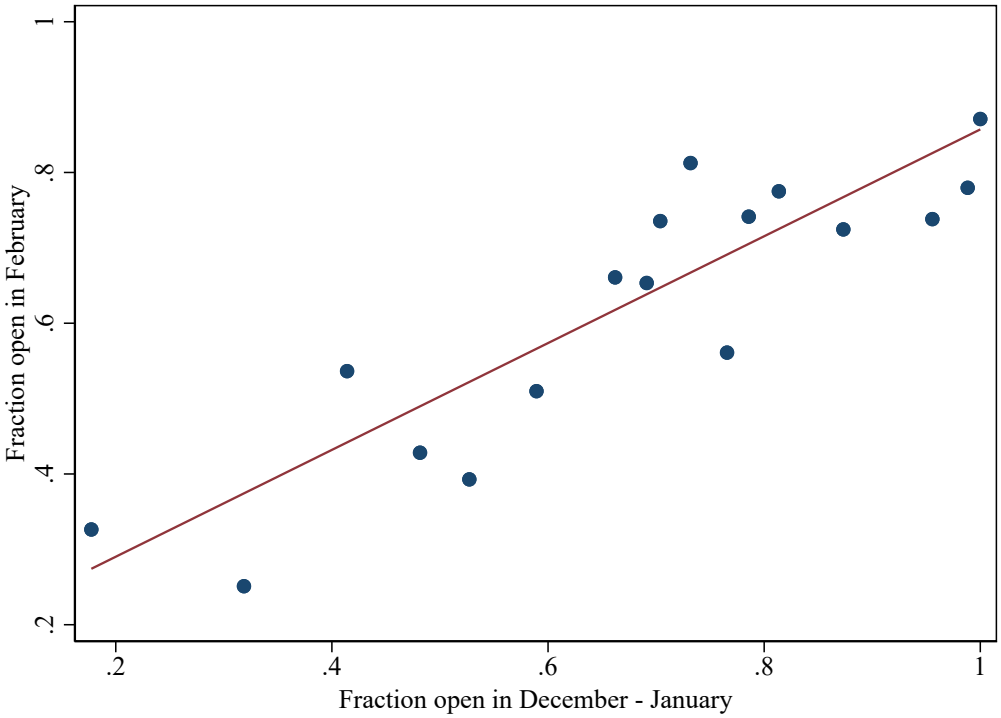
where $P_{i,t}$ is the price of a forward contract on a bulb (bulbs) of variety i traded in period t, τ_t are period fixed effects, $X'_{i,t}$ are control variables, specifically the log(weight) of bulbs before planting and dummies for the type of bulb traded (offset, single bulb, pound of bulbs, or multiple pounds of bulbs), and ψ_i are variety fixed effects. The sample is restricted to all varieties that have been traded at least twice. Periods t are defined as intervals wide enough to ensure sufficient transactions for estimation: $t \in$ {November 1-30, December 1-31 (1636), January 1-15, 16-27, 28, 30, 31, February 1, 2, 3, 4, 5, 5-9, May 1 (1637)}. Because sources sometimes record transactions over a multi-day span rather than a single date, we create 10,000 datasets in which such transactions are allocated randomly across the relevant days using a uniform distribution. Plotted coefficients are means across datasets. The 95% confidence intervals (truncated at 10 for readability) are constructed from the mean standard deviation across datasets, with standard errors clustered at the variety level. November 1636 is the omitted category, normalized to 1.

Figure 2: Distribution of fraction of days the canal was open (December & January)



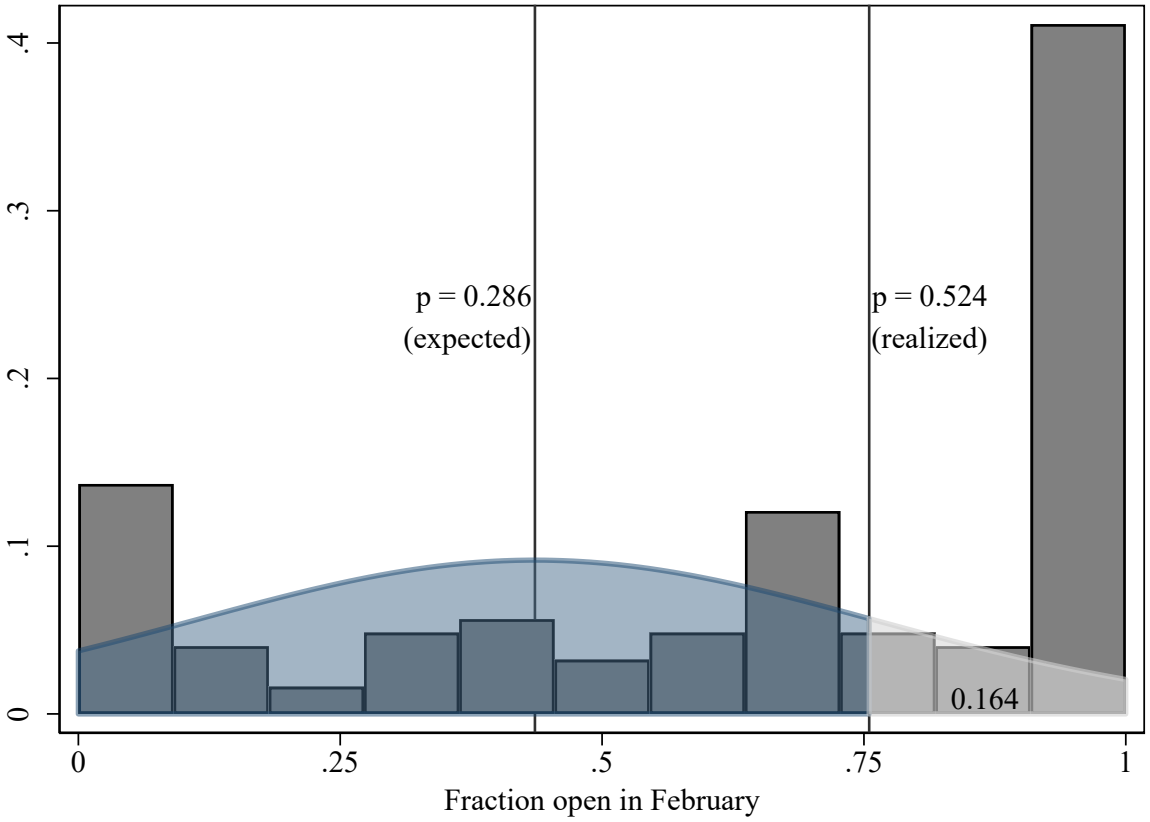
The histogram plots the distribution of the fraction of days in December and January that the canal between Amsterdam and Haarlem was open. Data are for December 1632 – December 1757. The vertical line indicates the observation for Dec 1636 – Jan 1637.

Figure 3: Fraction open: December – January vs February



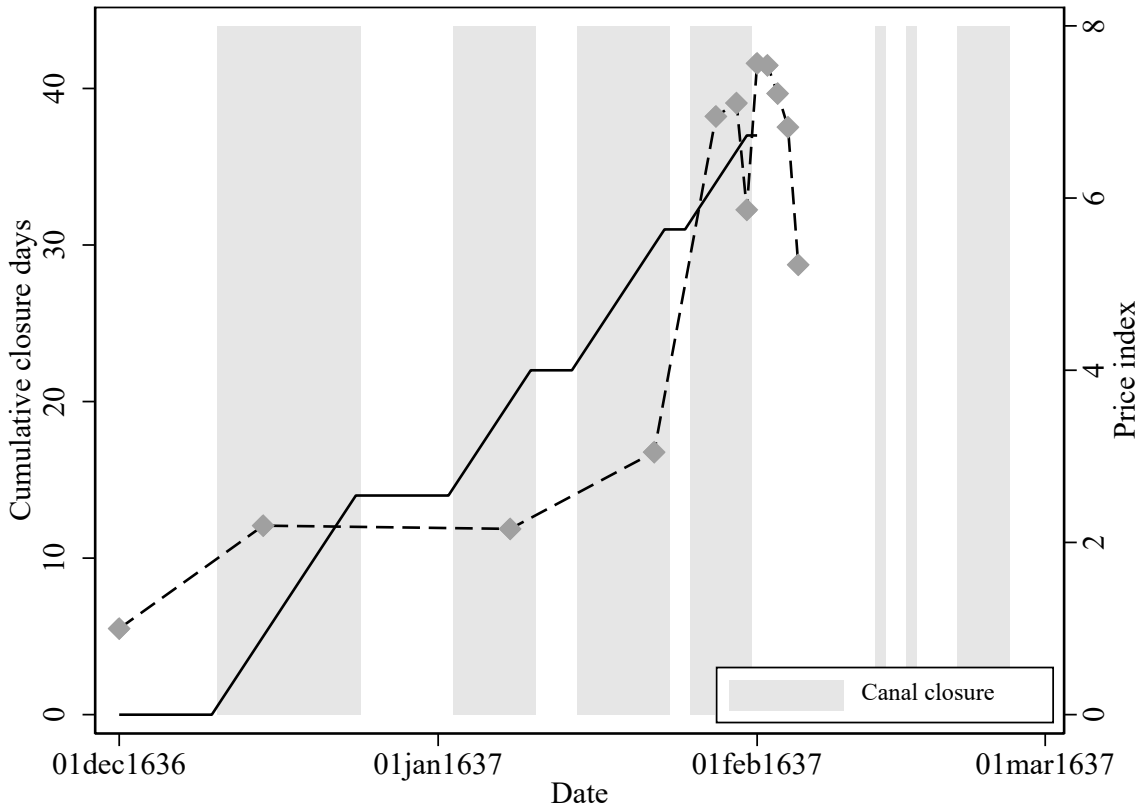
Binscatter of the fraction of days in December and January that the canal between Amsterdam and Haarlem was open against the fraction of days it was open in February. Data are for December 1632 – December 1757. Associated t-test and R2: 5.90 and 0.22.

Figure 4: Distribution of fraction of days the canal was open (February)



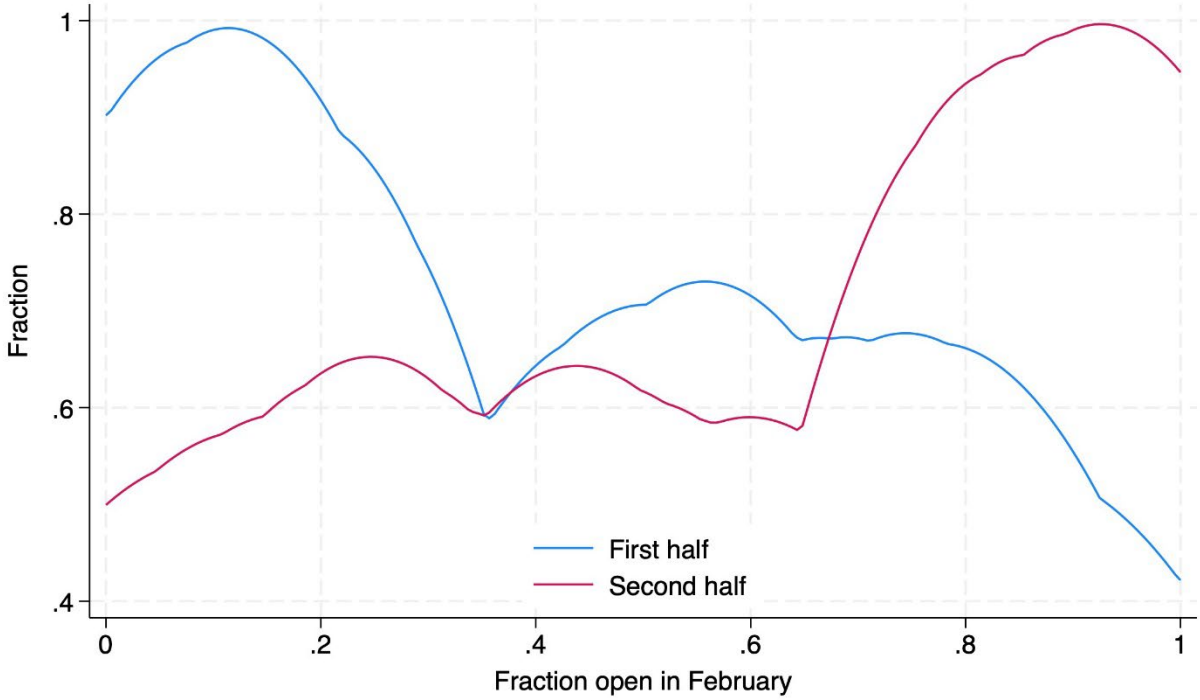
The histogram plots the distribution of the fraction of days in February that the canal between Amsterdam and Haarlem was open. Data are for December 1632 – December 1757. The two vertical lines indicate the fraction of days during February 1637 the canal was open, either realized or expected based on December and January, with their percentile position in the full distribution. The expected line is based on a linear regression and we report the estimated normal distribution around the point estimate taking into account both sampling variation and parameter uncertainty. The number in the lower right-hand corner gives the probability that the fraction the canal would be open is larger than the realization.

Figure 5: Days with closure and the estimate price index for forwards on tulip bulbs



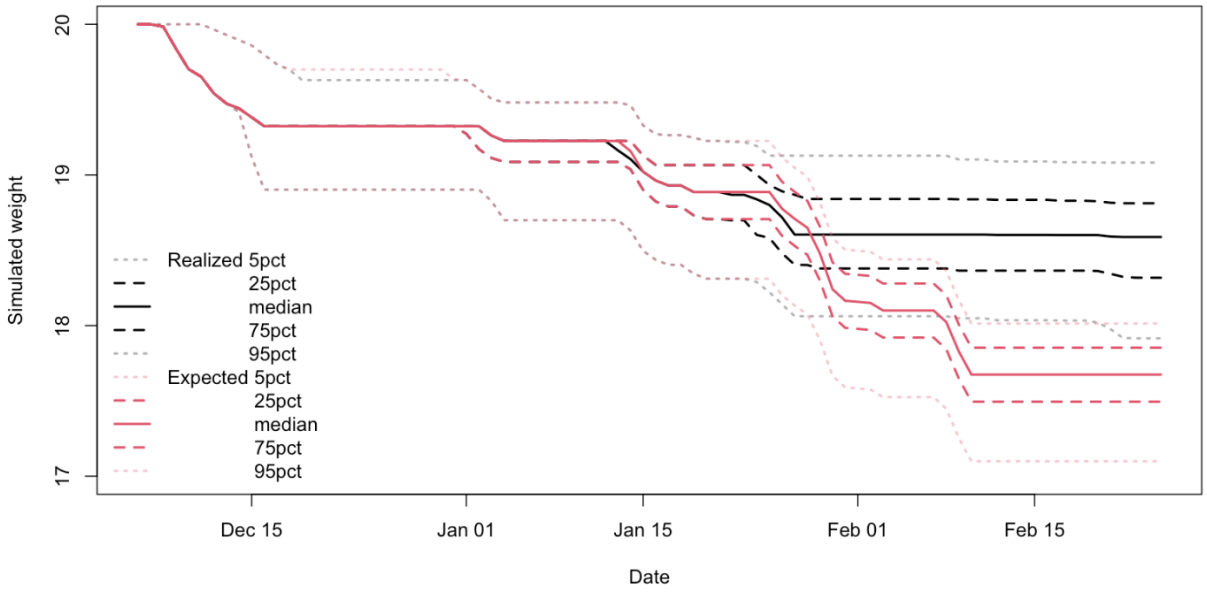
Daily and cumulative estimated closures of the canal between Amsterdam and Haarlem. Estimates are based on data from the Arnhem and Nijmegen canal, adjusted using the total number of days the canal between Amsterdam and Haarlem was actually closed. In particular, the Amsterdam-Haarlem canal is assumed closed when the Arnhem-Nijmegen canal reported "no traffic." Additional closure days are assumed to have taken place on days when the Arnhem-Nijmegen canal reported "little traffic" that immediately follow days with "no traffic." The estimated price index for forwards on tulip bulbs come from Figure 1.

Figure 6: Canal closures in February



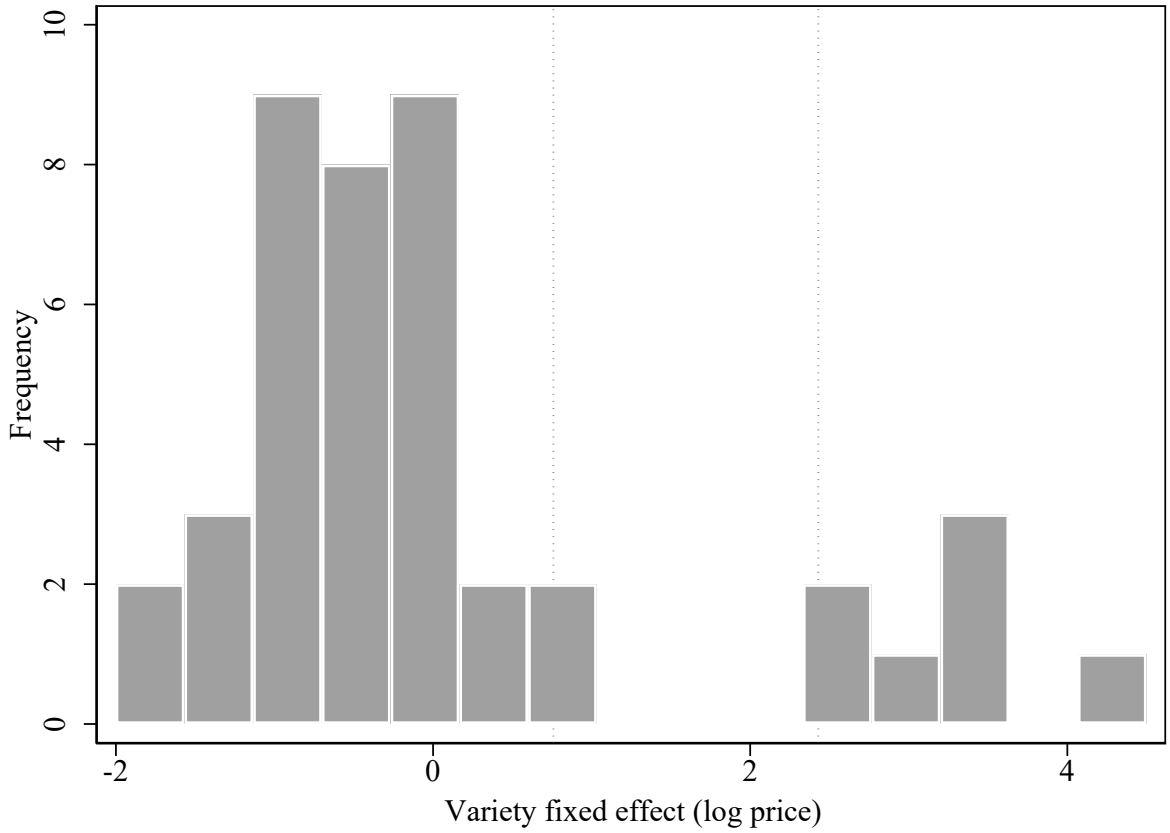
The figure shows kernel density plots (using the default bandwidth) of the fraction of days the canal between Amsterdam and Haarlem was open for the first and second halves of February. The underlying daily data are for 1632 –1757. The densities show that the canal was much more likely to be closed in the first half of February.

Figure 7: Estimated tulip bulb damage using expected and realized canal closures, 1636-1637



This figure shows the estimated damage to tulip bulbs during the winter of 1636–1637, based on simulated soil temperatures inferred from expected and realized canal closures. For expected canal closures, we assume that the expected number of closure days in February from Figure 4 occurred at the beginning of the month. We select ice segments from winters 1706–1756, for which daily temperature data are available, whose lengths closely match the realized and expected ice segments in 1636–1637 (±1 day for segments of ≤10 days, ±2 days for segments of >10 and ≤20 days, etc.). Using these segments, we generate 10,000 possible temperature series by randomly matching each realized or expected ice segment from 1636–1637 to a comparable segment drawn from the reference winters 1706–1756. When drawing segments, we weight matches by the similarity between the winter of 1636–1637 and the winter from which the segment is drawn. Similarity is measured by the fraction of days the canal was open over the entire winter; weights are defined as the inverse of the squared difference of this fraction and are normalized to sum to one for each realized or expected segment. For each of the 10,000 simulated temperature series, we compute the implied temperature-induced damage to a tulip bulb. The figure reports selected percentiles of the resulting damage distribution at each point in time across the simulations. Damage is defined as the weight of the sprout eventually growing out of the tulip bulb. Estimates are based on Van der Valk (1970) who provides experimental evidence on the relation between the length of exposure to below zero Celsius soil temperatures and the weight of a sprout, with a healthy sprout weighing 20 grams. The available data for 1706-1756 is for air temperatures and converted into soil temperatures using the restricted OLS estimates from Table A.1. in the Internet Appendix based on daily air and soil level temperatures during the winter of 1928-1929 (AVvB 1930).

Figure 8: Distribution of variety price fixed effects

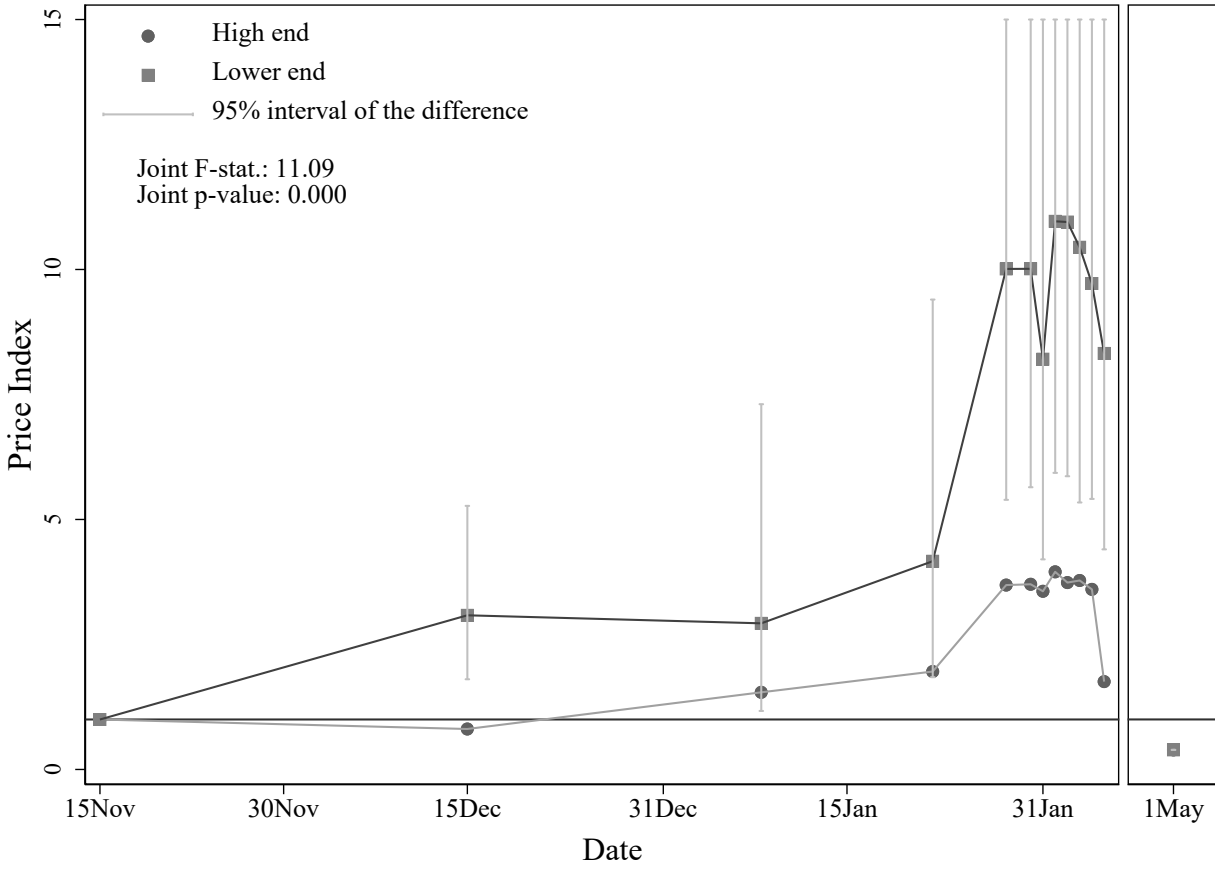


The graph reports histogram of the variety fixed effects (in logs) ψ_i estimated from:

$$\log(P_{i,t}) = \sum_{t \neq 15Nov} \beta_t \tau_t + X'_{i,t}\Gamma + \psi_i + \varepsilon_{i,t}$$

where $P_{i,t}$ is the price of a forward contract on a bulb (bulbs) of variety i traded in period t, τ_t are period fixed effects, and $X'_{i,t}$ are control variables, specifically the log(weight) of bulbs before planting and dummies for the type of bulb traded (offset, single bulb, pound of bulbs, or multiple pounds of bulbs). Periods t are defined as intervals wide enough to ensure sufficient transactions for estimation: $t \in$ {November 1-30, December 1-31 (1636), January 1-15, 16-27, 28, 30, 31, February 1, 2, 3, 4, 5-9, May 1 (1637)}. To estimate the variety fixed effects, we set the date of each transaction that is reported to have taken place over a certain interval to the midpoint of that interval. The sample is restricted to all varieties that have been traded at least twice; therefore, the variety fixed effects give the price level of a traded variety corrected for the general movement in prices. The two dotted vertical lines indicate the part of the distribution without any varieties, suggesting a clear differentiation between low and high-end varieties. The latter includes the Admiraels Liefkens, Van der Eyck, and Van Enckhuysen, and the Bruyne Purper, Generael der Generaelen van Gouda, Scipio, and Vice Roy.

Figure 9: Estimated price indices of high and lower-end varieties, 1636-1637



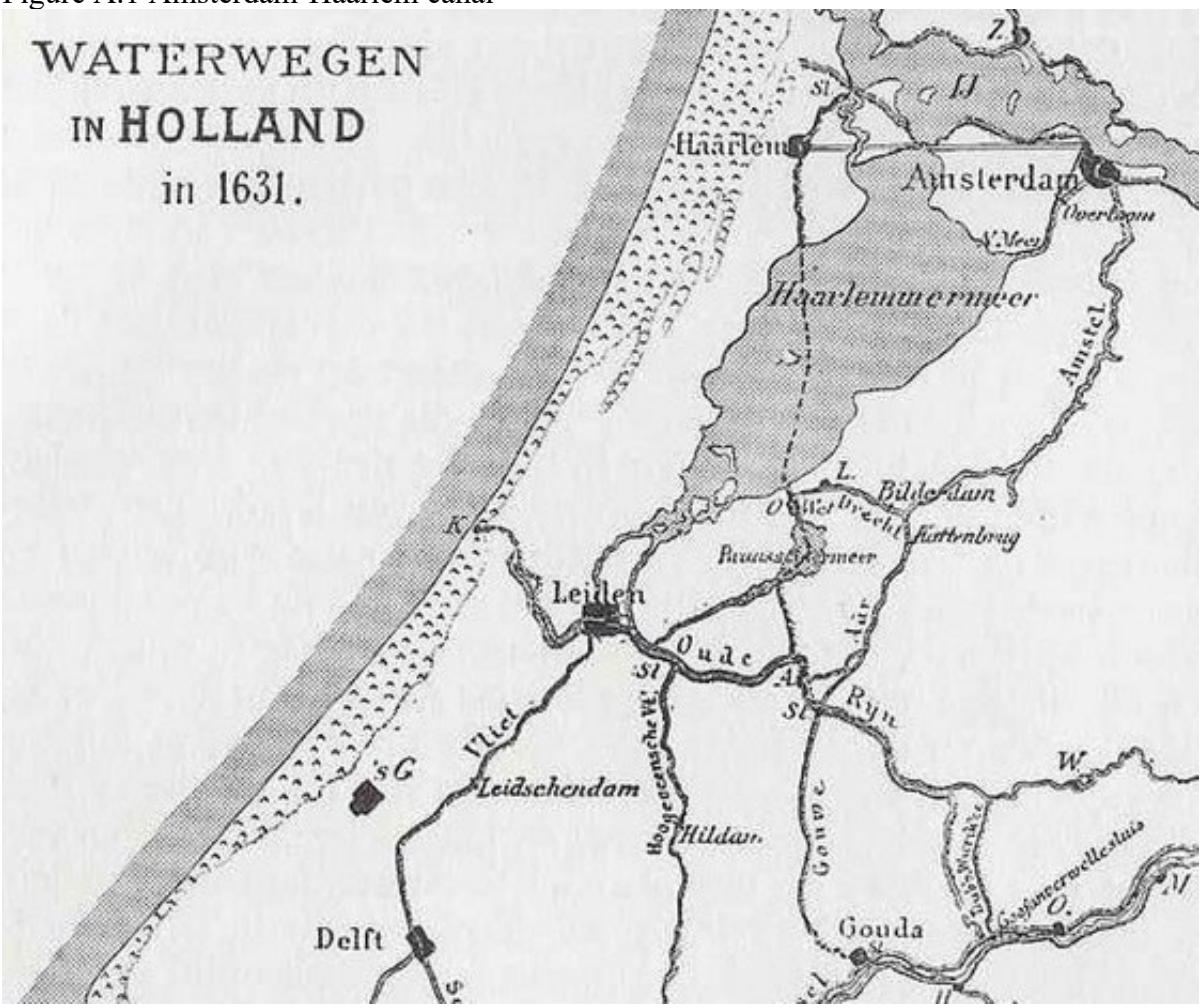
The graph reports the exponential of the time fixed effects from Figure 1 for high and lower-end tulip varieties separately. *High end*: defined in Figure 8; *Lower end*: all others. Specifically, we estimate:

$$\log(P_{i,t}) = \sum_{t \neq 15Nov} \beta_t \tau_t + \sum_{t \neq 15Nov} \gamma_t \tau_t \times L_i + X'_{i,t}\Gamma + \psi_i + \varepsilon_{i,t}$$

where $P_{i,t}$ is the price of a forward contract on a bulb (bulbs) of variety i traded in period t, τ_t are period fixed effects, L_i is a dummy for lower end varieties, $X'_{i,t}$ are control variables, specifically the log(weight) of bulbs before planting and dummies for the type of bulb traded (offset, single bulb, pound of bulbs, or multiple pounds of bulbs), and ψ_i are variety fixed effects. The sample is restricted to all varieties that have been traded at least twice. Periods t are defined as intervals wide enough to ensure sufficient transactions for estimation: $t \in$ {November 1-30, December 1-31 (1636), January 1-15, 16-27, 28, 30, 31, February 1, 2, 3, 4, 5-9, May 1 (1637)}. Because sources sometimes record transactions over a multi-day span rather than a single date, we create 10,000 datasets in which such transactions are allocated randomly over the relevant days using a uniform distribution. Plotted coefficients correspond to β_t and $\beta_t + \gamma_t$ and are means across datasets. The 95% confidence intervals (truncated at 15 for readability) are for the estimated coefficients γ_t centered on $\beta_t + \gamma_t$. As such they indicate the statistical significance of the difference between high and lower end varieties. They are constructed from the mean standard deviation across datasets, with standard errors clustered at the variety level. November 1636 is the omitted category, normalized to 1. The reported F-statistic and p-value come from testing the joint significance of the γ_t coefficients. There were no high-end tulip bulbs traded on May 1 and the figure only reports the time fixed effects for lower-end bulbs for this period.

Internet Appendix

Figure A.1 Amsterdam-Haarlem canal



Waterways in Holland in 1631. The horizontal line between Haarlem and Amsterdam depicts the canal of which the closures provide information about the weather conditions used in this paper.

Figure A.2. Canals and barges

Panel A: River landscape with sailing boats and a barge (Salomon van Ruysdael, 1660)



Panel B: Winter landscape with skaters (Adam van Breen, ~1608-1630)



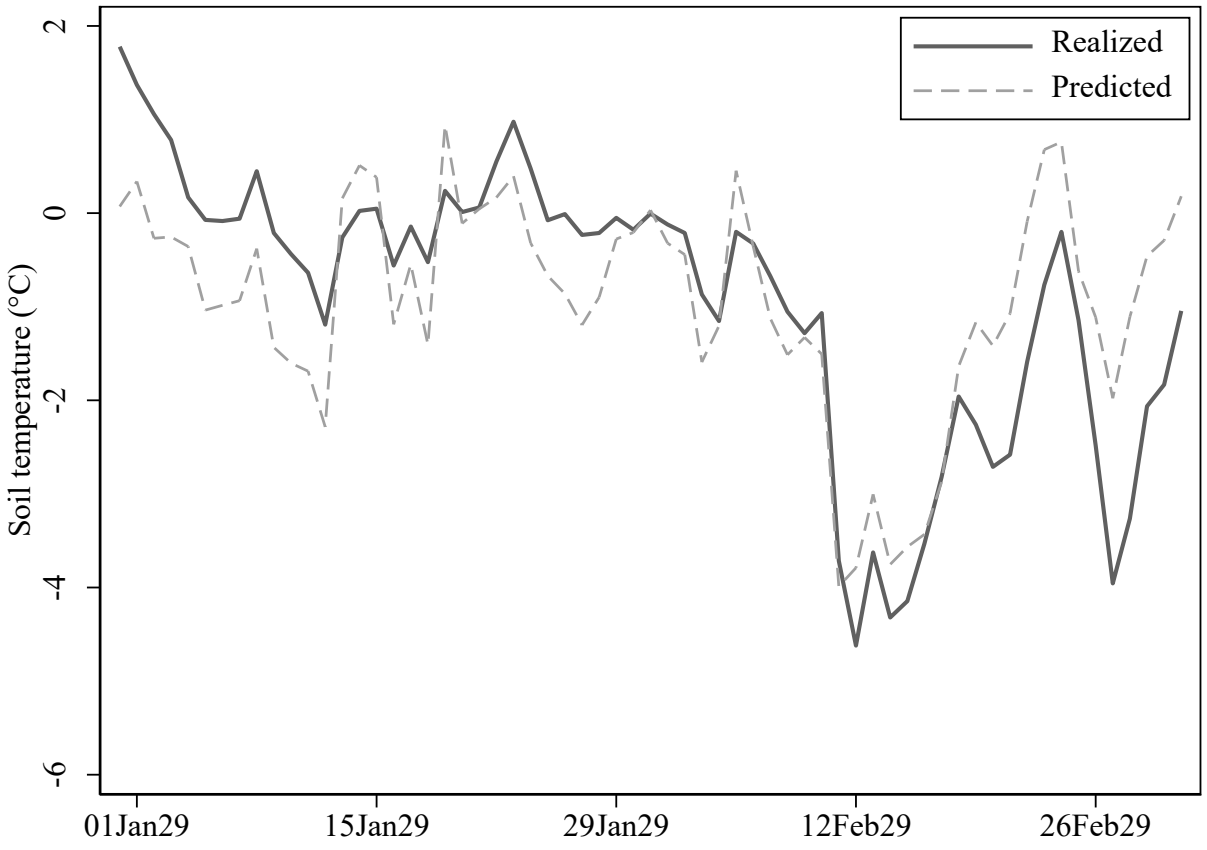
Panel A: view of a canal with a horse pulling a barge. Panel B: view of a frozen waterway and towpath normally used by horses to pull barges.

Table A.1: Predicting soil with air temperature

	(1) Soil temperature
Air temperature	0.304***
	(0.031)
Constant	0.123
	(0.158)
N	63
R^2	0.616
adj. R^2	0.609

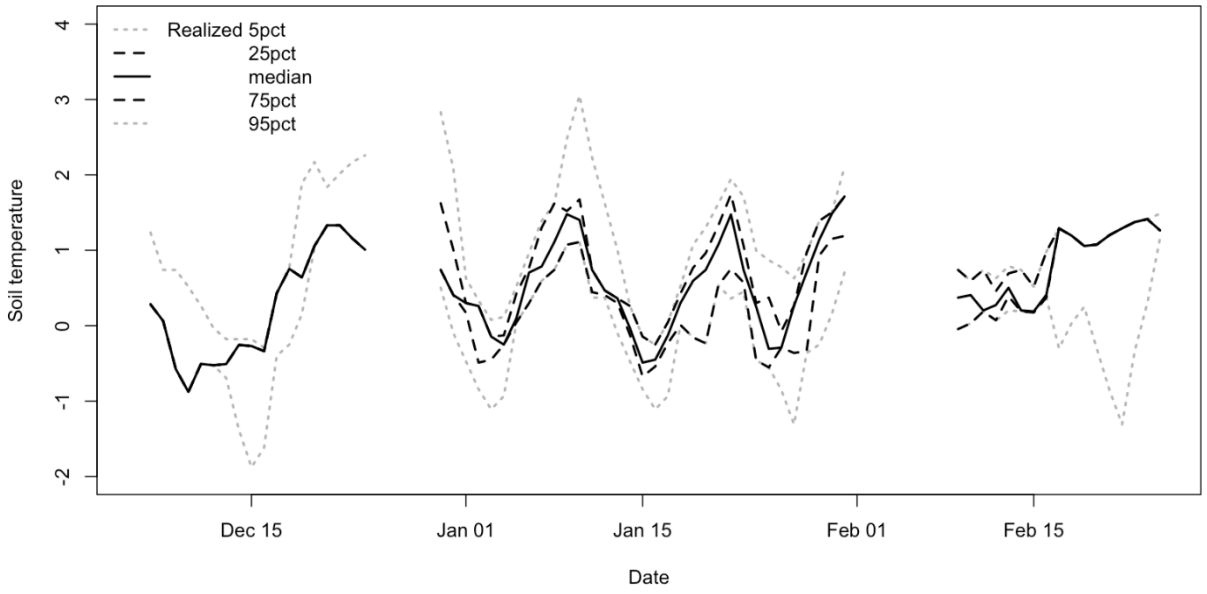
OLS estimates predicting today's soil temperature (at a depth of 6.5 cm with a cover of common reed) with today's air temperature. Period: January – February 1929. Both temperature series were measured at the experimental tulip bulb field in Lisse, about 15 km South of Haarlem. Source: AVvB (1930).

Figure A.3: Estimated soil temperature using air temperatures, 1928-1929



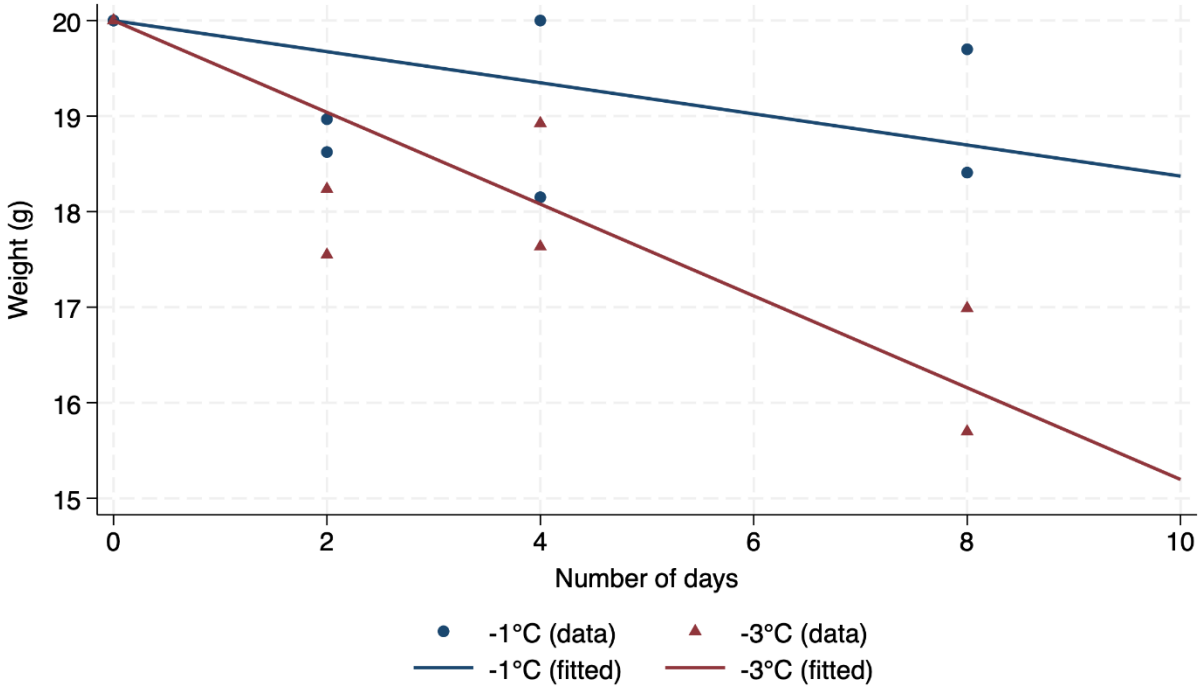
This graph show realized soil temperatures and those predicted by the OLS estimates of Table A.1.

Figure A.4: Estimated soil temperature using realized ice closures, 1636-1637



This figure shows the estimated soil level temperature during the winter of 1636–1637 using realized canal closures. We select ice segments from winters 1706–1756, for which daily temperature data are available, whose lengths closely match the realized segments in 1636–1637 (±1 day for segments of ≤10 days, ±2 days for segments of ≤20 days, etc.). Using these segments, we generate 10,000 possible temperature series by randomly matching each realized or expected ice segment from 1636–1637 to a comparable segment drawn from the reference winters 1706–1756. We take the temperatures during the period with ice and the four preceding days. When drawing segments, we weigh matches by the similarity between the winter of 1636–1637 and the winter from which the segment is drawn. Similarity is measured by the fraction of days the canal was open over the entire winter and weights are defined as the inverse of the squared difference in this fraction, normalized to sum to one for each realized or expected segment. The figure reports selected percentiles of the resulting temperature distribution at each point in time across the simulations. The available data for 1706-1756 is for air temperatures and converted into soil temperatures using the estimates from Table A.1. Missing values indicate periods without canal closures for which we cannot calculate temperatures.

Figure A.5: Mapping from subfreezing temperatures to weight loss of bulb shoots



This figure shows the experimental data from Van der Valk (1970), mapping exposure to subfreezing temperatures of −1°C or −3°C into weight loss of the shoots that the bulbs will produce in spring. Regression lines are based on restricted OLS with the constant set at 20 – the baseline.