

‘Those who know most’: insider trading in 18th c. Amsterdam*

Peter Koudijs[†]

November 29, 2012

Abstract

This paper employs a natural experiment from financial history to study the process by which private information is incorporated into prices. I look at the market for English securities in the Netherlands during the 1770s and 1780s. Anecdotal evidence suggests that English insiders traded actively on their private signals, both in London and in Amsterdam. I reconstruct the arrival dates of sailing boats that transmitted information from London to Amsterdam and I look at the movement of English security prices between the arrivals of boats. The evidence is consistent with a Kyle (1985) model in which insiders trade on their private signals in a strategic way and private information is only slowly revealed to the market as a whole. The speed of information revelation in Amsterdam crucially depended on how long insiders expected it would take for the private signal to be publicly revealed. The importance of private information is underlined by the response of London prices to price discovery in Amsterdam.

JEL Codes: G14, N2, D82, F36

Keywords: private information, strategic insider trading, financial history

*This paper is based on the second chapter of my dissertation. I would especially like to thank Hans-Joachim Voth (advisor), Fernando Broner and Jaume Ventura (thesis committee) for advice and guidance. I am grateful to seminar participants at Anderson-UCLA, UC Berkeley, Booth-Chicago, Columbia, Gerzensee, Harvard, Northwestern, NYU-Stern, UPF-CREI, Stanford, UC San Diego, Tilburg, WFA-Las Vegas, Warwick, Wharton-UPenn and Jonathan Berk, Jos van Bommel, Harry Garretsen, Vincent van Kervel (discussant), Floor Koudijs, Thijs Knaap, Larry Neal, Christine Parlour (discussant), Paul Pfleiderer, Giacomo Ponzetto, Juan Carlos Rodriguez (discussant), Christina Romer, David Romer, Dick Roll, Jonathan Rose, Subra Subrahmanyam (discussant), and Jan de Vries for their comments and suggestions. Many thanks to Daphne Acoca, Dmitry Orlov, Sophia Paliza and Hannah Rusk for research assistance. Any errors are my own.

[†]Stanford Graduate School of Business - koudijs@stanford.edu

Private information is central to our understanding of financial markets.¹ There is substantial evidence to suggest that it has a significant impact on asset prices. However, the process by which private information is incorporated into prices is difficult to analyze empirically. By definition, private information is not observed directly. It is not clear how insiders or informed agents trade on their information and by what process their private signals are revealed to the market as a whole.

Theory offers two perspectives. According to one view, agents with private information act in a competitive fashion. They do not take their impact on prices into consideration. They trade right after they receive a private signal and they do so aggressively. As a result, the privately-informed immediately reveal (most or all) of their private information.² According to the second view, insiders are strategic and take the price impact of their trades into account. They internalize that their profits fall as prices become more informative. This constrains their behavior. Trades are spread out over time. Price discovery - the process by which private information is incorporated into prices - is prolonged (Kyle 1985).³ Strategic behavior is not a sufficient condition for slow price discovery. Private signals also need to arrive relatively infrequently. If they arrive continuously, strategic insiders may find it optimal to reveal their current private information as quickly as possible (Gromb and Vayanos 1998; Caldentey and Stachetti 2010).

This paper uses a natural experiment from the 18th century to examine the incorporation of private information into prices in unique detail. Specifically, I test for the strategic behavior of insiders. I also study the speed of price discovery. The evidence supports the original Kyle (1985) model. It suggests that 18th century insiders traded in a strategic way, and that private information was only slowly incorporated into prices.

The evidence is based on a number of English securities that were traded in London and Amsterdam during the 18th century. Most, if not all, relevant information about these securities originated in London or reached London first and was then transmitted to Amsterdam (see section 1 for details). Anecdotal evidence strongly suggests that a significant fraction of this information was private in nature, and held by insiders (Sutherland 1952). I compare the price dynamics in the two markets to identify how this private information became incorporated into prices.

¹See for example Grossman and Stiglitz (1980); Kyle (1985); Glosten and Milgrom (1985); Admati and Pfleiderer (1988).

²If insiders are risk neutral and not capital constrained, their private information will be perfectly revealed (Holden and Subrahmanyam 1992, 1994). Vives (1995) analyzes the case where insiders are risk averse. He shows that, in the presence of a competitive market making sector, insiders will trade only once after they receive their information. This reveals their signal imperfectly (compare Grossman and Stiglitz 1980), but no additional private information is incorporated into prices afterwards. If there is no competitive market making sector, multiple equilibria exist (Grundy and McNichols 1989 and Brown and Jennings 1989). For more on competitive REE models see Wang (1994); He and Wang (1995); Llorente et al. (2002); Banerjee et al. (2009).

³In the classical Kyle model there is a single insider. Foster and Viswanathan (1996) show that the main predictions of the model are robust to the introduction of multiple insiders, as long as agents' private signals are (sufficiently) heterogeneous. Compare Holden and Subrahmanyam (1992, 1994).

The identification relies on the communication technology of the time. In the 18th century people relied on sailing ships to transmit news across the North Sea (Neal 1990; Koudijs 2012; see figure 2 for a map). There was an official mail packet boat service that carried both public newspapers and private letters. For all practical purposes this was the only way in which public and private information got transmitted from London to Amsterdam.⁴ On average these sailing boats arrived twice a week, but due to adverse wind conditions at sea this could vary considerably.

I study the process of price discovery in the following way. Under the null hypothesis all public information was immediately incorporated into prices. If there was private information as well, this was held by competitive insiders. Depending on whether these insiders were risk neutral or risk averse this information was either fully or partly incorporated into prices. In both cases insiders traded immediately on the information and prices moved only once in response (see footnote 2 on page 2). If the Amsterdam market responded to this information efficiently, Amsterdam prices immediately reflected the London information after a boat arrival. Any other price movements in Amsterdam between the arrival of two boats should be unrelated to developments in London.

Under the alternative hypothesis private information in London was held by strategic agents. While public information was immediately incorporated into prices, private information was not. London insiders traded strategically and price discovery was spread out over time. The packet boats transported both public and private information. Amsterdam insiders observed the private information and, if they also traded strategically, price discovery in Amsterdam was spread out over time as well.⁵ As a consequence, prices in Amsterdam not only responded to London on the days the packet boats came in. More generally they moved in the same direction as prices in London, reflecting the incorporation of the same private signals. This intuition is illustrated by figure 1.

[FIGURE 1 ABOUT HERE]

The data rejects the null hypothesis in favor of the alternative. The co-movement patterns of prices in London and Amsterdam are consistent with slow price discovery and supports the original Kyle (1985) model in which privately informed agents trade strategically. There are two additional pieces of evidence that corroborate this conclusion. First of all, a Kyle-type model predicts that if insiders expect to have a lot of time to use their private signal, they will trade less aggressively early on. This ensures that their information is not revealed too quickly and they continue to make insider profits in later periods (for a formalization of this intuition see section 2). This is confirmed by the data. Based on the sailing schedule and local wind conditions I reconstruct when agents in Amsterdam expected the next boat to arrive. This boat carried information from London (e.g. prices) that (partly) revealed the insiders' signals. Consistent with the theory, the initial co-movement of Amsterdam and London prices was weaker when the next boat was expected to arrive late and there were more opportunities to trade on his signal.

⁴The same packet boat service also brought Amsterdam news back to London. See section 1 for details.

⁵The Amsterdam insiders could either be agents or business partners of the London insiders, see section 1 for examples.

Second, if prices in Amsterdam reflected the incorporation of private information, then the London market could have learned from Amsterdam prices.⁶ This is confirmed by the data. Conditional on its own price discovery, the London market updated its beliefs based on price changes in Amsterdam. If due to weather conditions it took longer for a private signal to "bounce off" from Amsterdam, the London market responded less. When a roundtrip London-Amsterdam took longer than two weeks, feedback from Amsterdam became irrelevant. At that point the initial private signal was already fully incorporated into London prices and the Amsterdam signal became irrelevant.

I test for a number of alternative explanations. First of all, the co-movement of Amsterdam and London prices might be explained by public information that was only slowly incorporated into prices. This could be due to significant trading costs (or limits to arbitrage) that prevented prices from immediately adjusting to news. In one interpretation, this would lead to momentum in the return series (see Hong and Stein 2001). Although there is some evidence for return continuation, this does not drive the results.⁷ In another interpretation prices may not adjust at all because of trading costs (Lesmond et al. 1999; Bekaert et al. 2007). Price adjustment only takes place when the news shock is big enough or when trading costs are low. I show that co-movement between Amsterdam and London was slightly stronger after zero returns in London. However, this does not drive the results.

Secondly, the co-movement between Amsterdam and London could be driven by correlated liquidity shocks rather than by private information. As liquidity traders move down (potentially) downward sloping demand curves in both markets, prices move in the same direction. It is natural to assume that liquidity shocks have a transitory impact on prices (Grossman and Miller 1988). In other words, liquidity shocks should lead to return reversals. If this drove the co-movement between Amsterdam and London, then we should observe return reversals across markets. I.e. positive (negative) price changes in London should predict subsequent negative (positive) returns in Amsterdam. This is not the case.

Finally, co-movement might simply be driven by news slipping through other channels than the official packet boat service. There is no qualitative evidence that market participants relied on sources other than the official mail service. The most important Anglo-Dutch bank of the period that was active in the stock market (Hope & Co) fully relied on the official mail. Nevertheless it is possible that other market participants used alternative ways to transmit information, specifically other sailing boats. To test for this possibility I restrict the sample to periods where, after the arrival of a packet boat, wind conditions suddenly turned so that future packet boats were significantly delayed. I assume that during these periods it was equally impossible for other boats to sail across

⁶This is true as long as noise trading shocks were not perfectly correlated across markets. If this condition is met, price discovery in two markets is more informative than in one (Chowdry and Nanda 1991; Boulatov et al. 2011).

⁷Note that the discovery of private information in a Kyle model (or any type of REE model) does not lead to momentum. The intuition behind this is that in a REE, future price changes should not be predictable based on past returns (see section 2 and Banerjee et al. 2009).

the North Sea. If co-movement was purely driven by news arriving through other channels, prices should not move in the same direction in London and Amsterdam during these episodes. I find that co-movement was present and just as strong during periods of adverse wind conditions. This does not disprove that channels other than the packet boats were used. However it does indicate that private information is needed to fully explain the co-movement patterns in the data.

Note that none of these three alternative explanations can explain why co-movement was stronger (weaker) if the next boat was expected to arrive early (late). Nor can they account for the conditional feedback effect of Amsterdam price changes on London. Without a Kyle model, it is difficult to make sense of these findings.

This paper is related to two strands in the existing literature. First of all, there is a large body of empirical literature that documents the importance of private information for asset price movements. Most empirical work interprets the price impact of transactions (or order flow imbalance) as evidence for the relevance of private information (Hasbrouck 1991; Easley, Kiefer and O'Hara 1997; Madhavan et al. 1997; Evans and Lyons 2002; Vega 2006 and related papers). Recently this interpretation has attracted some criticism because the price impact of order flow imbalance can capture liquidity as well as information (Duarte and Young 2009). There have been alternative approaches to study the impact of private information. For example, Pasquariella and Vega (2007) and Tetlock (2010) look at the interaction of order flow imbalance (or volume) and public information events. Marin and Olivier (2008) find that stock prices drop sharply after (reported) insider sales peak. Cohen et al. (2008; 2010) focus on the performance of institutional investors in stocks of companies that are run by former class mates. Cohen et al. (2011) analyze reported insider trades and show that non-predictable trades outperform the market. Kelly and Ljungqvist (2011) look at the impact of the closure of brokerage firms' research departments. The present paper complements these findings and provides further evidence for the relevance of private information. The key strength of this paper is that 18th century financial markets were less complex than today and that information flows can be perfectly reconstructed. This allows for a clean identification of private information.⁸

Most importantly, this paper is related to the questions whether insiders are competitive or strategic and, related to this, how quickly private information is incorporated into prices (for theoretical contributions see for example Kyle 1985; Glosten and Milgrom 1985; Subrahmanyam 1991; Back 1992; Holden and Subrahmanyam 1992, 1994; Foster and Viswanathan 1994, 1996; Romer 1993; Chau and Vayanos 2008; Caldentey and Stacchetti 2010). Because private information is by definition unobservable, there is only limited empirical evidence available to answer these questions. Most evidence is indirect and based on intra-daily volatility and volume patterns (see inter alia Madhavan et al. 1997; Dahya et al. 2010). Notable exceptions are Boulatov et al. (2011) and Hendershott et

⁸In related work (Koudijs 2012) I estimate what fraction of overall return volatility in Amsterdam can be attributed to the arrival of public and private information. In that paper I do not delve into how price discovery exactly worked – the key question addressed in this paper.

al. (2012) who show that the trading behavior of institutional investors is consistent with a Kyle model. However, it is not clear to what extent institutional order flow can capture informed trading.

Relative to this literature I make the following contributions. First of all, the paper provides direct evidence that insiders acted strategically and traded differently depending on the time they expected to have to benefit from their private information. This significantly affected the speed of price discovery (Caldentey and Stacchetti 2010). Secondly, the paper finds that private information is slowly incorporated into prices, confirming the predictions of Kyle (1985) and related contributions. I estimate that it took approximately two weeks for a given private signal to be incorporated into prices.

The rest of the paper is organized as follows. Section 1 discusses the historical background and context of this paper in more detail. In addition, I provide anecdotal evidence for the relevance of private information. To motivate the empirical analysis I set up a simple Kyle model in section 2. Section 3 presents empirical evidence that supports the model's predictions. Section 4 provides a number of robustness checks and extensions. Section 5 concludes.

1 Historical background

In separate work (Koudijs 2012, see also Neal 1990) I provide a more detailed overview of the market for English stocks in Amsterdam in the 1770s and 1780s. In this section I summarize this historical background and I give ample attention to the microstructure of this market.

1.1 Stock and sample period

The data necessary for this paper's analysis are available for three different English securities, British East India Company (EIC) stock, Bank of England (BoE) stock and a government bond, the 3% Annuities.⁹ The empirical analysis in the main text focuses on the EIC. Results for the other two securities, which are overall very similar, are presented in Appendix B. The EIC was a trading company that held large possessions in today's India. The company's prospects were to a large extent determined by conditions in India. However, during the second half of the 18th century political developments in England started to become of key importance.¹⁰ (Sutherland 1952). The BoE operated to help finance the British government debt. The BoE was set up in 1694 to function as the government's banker. In addition, the BoE also provided large scale credit to the EIC. Finally, it discounted commercial bills, but on a relatively modest scale (Clapham 1944).

⁹In addition to these three securities, South Sea Company stock and 4% Annuities were also widely traded in Amsterdam. However, there is no frequent price data for London available for these securities.

¹⁰There was a constant discussion inside and outside the British Parliament about the semi-private character of the company and its public function. In addition, the company required regular bailouts from the English government to stay on its feet. As a result, political gyrations had an important impact on the company's share price

The analysis of this paper rests on the assumption that all relevant information about the English securities was generated in England. This is not necessarily true for the entire 18th century (see Dempster et al. 2000). The period was filled with European continental wars or the threat of a war breaking out, and England was involved in nearly all of them (Neal 1990). I therefore limit the empirical analysis to the sample periods 1771-1777 and 1783-1787. Both periods are characterized by peace on the European continent.¹¹

In Koudijs (2012) I show that during these two periods most relevant information about the English securities originated in England. Figure 14 in Appendix B presents the impulse response functions of price changes of EIC stock in Amsterdam responding to London and vice versa. The figure shows that Amsterdam responded strongly to London, with hardly anything going in the other direction.

1.2 The flow of information between London and Amsterdam

How exactly did English news reach Amsterdam? England and the Dutch Republic were connected through a system of sailing ships, at the time referred to as packet boats. The system was organized between Harwich and Hellevoetsluys, a important harbor close to Rotterdam (see figure 2). Since Amsterdam did not have a direct connection with the North Sea, this was the fastest way information from London could reach Amsterdam (Hemmeon 1912; Ten Brink 1957, 1969; Hogesteeger 1989; OSA 2599).

[FIGURE 2 ABOUT HERE]

Each packet boat brought in newspapers and other public newsletters with information about the recent developments in London, including the most recent stock prices. In addition, the packet boats brought in private letters. These could be simple letters from London correspondents with political and economic news and updates about stock market conditions¹². They could also be private letters from London insiders to their agents in Amsterdam; the focus of this paper.

The packet boats were scheduled to leave on fixed days: Wednesday and Saturday. The median sailing time was 2 days (including the day of departure). It took additional time to transport the news over land. Roads were particularly bad during the period and rivers had to be crossed by

¹¹The starting point of the first period, September 1771, is determined by data limitations. The period stops in December 1777 as tensions between France and England increased, eventually leading to outright naval war in July 1778. The second sample period starts in September 1783, right after the signing of an official peace treaty between France and England. There had been an armistice between France and England since January 1783 and the official peace treaty meant a return to normality. The second sample period stops in March 1787 when domestic tensions in the Netherlands rose, eventually leading to minor skirmishes in May 1787 and an intervention by the Prussian army in September 1787.

¹²Wilson (1941, pp. 74-75) gives a number of examples where people with an interest in the English stocks received private correspondence from London. For other examples of such letters see the correspondences the Amsterdam broker Robert Hennebo and the bankers Hope & Co maintained with their agents in London (Van Nierop 1931, *passim* and SAA 734; 78,79, 115 and 1510) and the estate of the Jewish broker Abraham Uziel Cardozo (SAA 334; 643).

ferry. Even though the news was transported on horseback, this still took considerable time, adding two days, making a total of 4 days (including the day of departure).¹³

The sailing ships often encountered adverse winds and as a result the news could be significantly delayed for days, sometimes even weeks. Around a third of the North Sea crossings from England to Holland were delayed. The longest delay I encountered was 17 days. As a result there was considerable variation in the time between the arrival of boats. During these periods of bad weather no news was transmitted across the North Sea. The *Tatler*, an English newspaper of the time, described that there could be a news blackout in London "when a West wind blows for a fortnight, keeping news on the other side of the Channel" (Dale 2004, p. 17). The same was true for Amsterdam when the wind was blowing from the East.

The packet boat system was the main source of English information for investors in Amsterdam, including insiders. The Dutch newspapers of the time all relied on the packet boat service to get news from England (*Amsterdamsche Courant*; *Opregte Haerlemsche Courant*; *Rotterdamsche Courant*). During the sample period, all articles in the *Amsterdamsche Courant* with new information from London can be retraced to the arrival of a specific packet boat, except for a number of exceptions I discuss below.

Even Hope & Co., one of the biggest Dutch banks of the period, with strong connections in London and heavily engaged in insider trading in Amsterdam (see p. 10) seems to have relied solely on the packet boat system to receive its private dispatches from London.¹⁴

At times, during periods of particularly bad weather, the English news could arrive in Amsterdam through the harbor of Ostend in today's Belgium, which had a regular packet boat service with Dover in England. During bad weather episodes it was impossible for the packet boats to sail between Harwich and Hellevoetsluys but other mail boats seem to have managed to get across to Ostend. With a total of nine times this happened only infrequently during the entire sample period. These episodes were meticulously reported by the Dutch newspapers and I account for them in the empirical analysis.¹⁵

The packet boats were of course not the only ships that sailed between London and Amsterdam.

¹³In London news would be collected by the end of the day on Tuesday or Friday (day of departure: day 1). This was transported to Harwich in the early morning, from where a mail packet boat would set sail in the afternoon (day 2). The boat would usually arrive in Hellevoetsluys on the next day (day 3). After the news had arrived it was quickly sent to Amsterdam where it usually arrived the day after (day 4).

¹⁴Most English letters in the Hope archive mention both the date a letter was written in London and the date it was received and opened in Amsterdam. I found 112 letters that Hope received from London during the sample periods. 99 of these letters were dated on mail days and were written right before the next mail boat would leave. For 83 letters, I could identify on what day Hope received and opened these letters. Out of these 83 letters, 73 were received on days the mail packet arrived in Amsterdam. Five letters were opened one day late when the news had arrived in the evening of the previous day. The final five letters were for some reason only opened a number of days later (Hope & Co, SAA 735: 78,79, 115 and 1510).

¹⁵I did not find a single reference to English letters received over Calais. Apparently, from a Dutch perspective, the Ostend connection always beat the Calais one.

Each week ships coming from England would dock in the Amsterdam harbor. However in terms of keeping up with current affairs these ships were always behind the packet boats. Amsterdam had no direct connection to the North Sea and boats had to sail via the isle of Texel to enter Amsterdam from the east. This would take a number of additional days. It therefore comes as no surprise that both individuals and the public newspapers had to rely on the packet boat service to get the most recent news from London.

Although the packet boat service seems to have been the most important source of information for Dutch investors, the flow of news through alternative channels can never be completely ruled out. It is possible that investors set up private initiatives to get information from London. For example, market participants may have used carrier pigeons to get information from London. The use of pigeons can be retraced to antiquity. However the historical record suggests that people only started using them intensively after 1800 (Levi 1977). In Koudijs (2012, appendix D) I find no evidence for their use in 1770s and 1780s. Most importantly, carrier pigeons could not be used in the winter months and I find no differences in price patterns between the winter and the rest of the year.¹⁶

A more relevant possibility could be that investors hired private boats to transmit information from London to Amsterdam. For example, there are rumors from the South Sea bubble in 1720 that Dutch investors chartered their own fishing ships to get the most recent information from London (Smith 1919; compare Jansen 1946 who found no evidence for these rumors). It seems reasonable to assume that if the official packet boats could not sail because of adverse weather conditions it would have been extremely difficult for other boats to cross the North Sea. In section 4.1 I use this logic to perform a number of robustness checks on the main results.

1.3 Market microstructure

During the 18th century the stock trade in Amsterdam took place in a decentralized fashion. Around noon there were two official trading hours in front of the Exchange building (Spooner 1983; Hoes 1986). However, trade continued outside these official hours in coffee shops and even in front of the Jewish synagogue (many traders were Jewish). Trading seems to have continued into the evening. A central clearing mechanism for the stock trade was missing and most trades took place through the direct matching of buying and selling parties (Van Nierop 1931).

This matching was done by a relatively small group of brokerage firms. Smith (1919) argues that in 1764, 41 brokerage firms were dominating the market. The correspondence of broker Robert Hennebo published in Van Nierop (1931) indicates that the market was driven by limit orders. Principals would transmit these orders to their brokers, who then tried to execute these orders to the best of their ability. This has an interesting implication for the prices we observe. The prices that

¹⁶There is a famous anecdote from 1815 where Nathan Rothschild received news about the outcome of the battle of Waterloo before anybody else by ways of a carrier pigeon. This story has been debunked as a myth (Ferguson 1998). His courier used a boat.

were reported most likely reflected the equilibrium price at which most limit orders could be cleared. The correspondence in Van Nierop (1931) suggests that prices were indeed interpreted this way.

By the second half of the 18th century a significant fraction of trade in the English stocks in Amsterdam was concentrated in the futures market. This has the important implication that it was relatively easy for market participants to go short as well as long¹⁷.

1.4 Private information

There is ample anecdotal evidence that London insiders used the Amsterdam market extensively to benefit from their private information. Insider trading in London and Amsterdam wasn't banned until the 1930's and especially EIC stock featured frequent insider trading. In a letter to one of his clients from January 1731 Amsterdam broker Robert Hennebo mentioned that there had been some active buying of EIC stock on the exchange and that

‘if I am not mistaken, these orders came from London, from one of the directors of the EIC, John Bance (...), making it likely that the share price will rise some more’.¹⁸

There is also evidence that EIC directors James Cockburn and George Colebrooke were ‘bulling’ the Amsterdam market during 1772 (Sutherland 1952, p. 228; SAA Hope, Journal 1772). One of his contemporaries would later describe Colebrooke as he, ‘who was in the secret, knowing when to sell for his own advantage’ (quoted in Sutherland 1952, p. 234). Such practices were not restricted to directors of the EIC. At times, political developments had a profound impact on the Company's prospects and as a result British politicians would engage in insider trading as well. During the 1760's a group of parliamentarians, amongst whom Lord Shelburne, a later prime-minister, and Lord Verney, member of the Privy Council, regularly engaged in insider trading in EIC stock. The Dutch banker Gerrit Blaauw traded for their account in the Amsterdam market (Sutherland 1952, pp. 206-8).

The clearest example of informed trading in Amsterdam comes from the archives of Hope & Co. In the fall of 1772 Hope went into business with Thomas Walpole to speculate on EIC stock. Walpole was a London banker but also the nephew of former Prime-Minister Robert Walpole, and a prominent Member of Parliament (Sutherland 1952, pp. 101 and 109). Walpole was clearly a political insider. On the 22nd of December 1772 Hope received a letter from Walpole dated the 18th which was labeled ‘private’ and read:

‘A report is made [on the poor state of the EIC] and we shall soon judge of its effect upon the stock. Those who know most think the stock will fall and we are of that opinion.

You may therefore resume your sales to such extent as you think proper and with the

¹⁷See the example on p. 11.

¹⁸Hier is gisteravont veel premy voor de reysing gegeven, en vandaag waren hier koopers tot 169. So ik niet mis heb komt die order van London, van een der directeurs, Mr. John Bance, sodat, gelyk ik Ued meermaals geseget en geschreeven heb, de apparentie grooter voor een reysing dan voor een daling is'. Van Nierop (1931), p. 68.

usual dex[t]erity. (...) There appears no risk in selling from 170% to about 166% [prices in % of face or nominal value]. One wouldn't go lower, for though it is probable the stock will fall to 150%, yet at that price or higher people may begin to speculate for the rise which will undoubtedly take place when any plan shall be fixed for the relief of the company. Whenever therefore the price falls to 154% or thereabouts, we should not only settle our positions but purchase more with a view to the rise as circumstances may make it advisable'.¹⁹

Walpole's intelligence proved to be accurate. On January 14, 1773 the Directors of the EIC asked for a government loan and concessions to export tea to all British colonies, both of which were granted (Sutherland 1952, pp. 249-251). Most importantly, Walpole's prediction on the price trajectory largely came true. The price of EIC stock in Amsterdam fell from 169.5% to 161%²⁰ on December 30, reaching its lowest point on January 4, 1773 at 157.50%²¹. After that, the price of EIC rose back to 169.5% on January 29, 1773.

Another interesting point is the reference to 'the usual dex[t]erity' Hope had to apply when executing the transactions. Most likely Hope had to be careful not to trade too conspicuously and reveal the information to other market participants. Hope probably did this by going through intermediaries. Hope's bookkeeping indicates that all share transactions on the Amsterstam exchange were handled by the brokerage firm David Pereira and Sons (SAA 734).

Unfortunately it is not possible to exactly reconstruct the profits Hope and Walpole managed to make by trading in Amsterdam. However, there are quarterly profit and loss statements available that indicate that on February 15, 1773 Hope credited Walpole £679:7:6 for profits on a short position in EIC stock of £18,500 nominal in the Amsterdam futures market (SAA 734). This short position must have run somewhere between December 15, 1772 and February 15, 1773. Assuming that Hope & Walpole shared the proceeds of this transaction 50/50, the profit from this short position is consistent with a price fall of -7.3 %-point (so for example from 165% to 157.7%).

1.5 Data

The empirical analysis of this paper is based on detailed price data from the Amsterdam and London markets and information about the arrival of packet boats in Hellevoetsluys (and Harwich). Data on British stock prices in Amsterdam were hand collected from the *Amsterdamsche Courant* and where necessary supplemented by the *Opregte Haerlemsche Courant*. Three prices a week are available for Monday, Wednesday and Friday.²² The Amsterdam market traded English stocks in Pounds

¹⁹Private letter from Walpole to Hope, Hope & Co., SAA 735; 115.

²⁰In the 18th century prices were generally reported in % of nominal/face value.

²¹In London the EIC stock price did fall to 154 on January 16th.

²²Previous research by Neal (1990) and Dempster et al. (2000) use Amsterdam prices with a frequency of 2 observations a month.

Sterling and prices were therefore reported in Pounds (as the percentage of nominal or face value). The Amsterdam prices the Dutch papers published were supplied at the end of the afternoon by a committee of so-called sworn brokers who were officially responsible for reporting these prices (Smith 1919, p. 109; Polak (1924); Jonker 1996, p. 147). Price data from London are available on a daily basis (Monday - Saturday) and are taken from Neal (1990), where necessary supplemented with Rogers (1902).

By the second half of the 18th century, a significant fraction of trade in the English stocks in Amsterdam was concentrated in the futures market (Van Dillen 1931). As a result all available price data for the Amsterdam market refers to futures prices²³. Prices in London are spot. In order to facilitate comparison between London and Amsterdam prices, I converted Amsterdam future prices into spot prices (for details see Koudijs 2012, Appendix A).

The arrival dates of boats in Hellevoetsluys were hand collected from the *Rotterdamsche Courant*. The newspaper reports on what day a specific boat arrived and whether it arrived before or after 12 p.m. This data can be used to determine when news from England must have arrived in Amsterdam. It took approximately 16 hours for news from Hellevoetsluys to be transported to Amsterdam (Stitt Dibden 1965, p. 9). This generally means that the information brought in on a certain day was only available to Amsterdam investors during the next day.²⁴ The *Rotterdamsche Courant* not only mentions the day a specific boat arrived but also the date of the news it carried. This information can be used to reconstruct what London price information was available to Amsterdam investors at certain points in time. In addition, I hand collected information about the arrival of Amsterdam news in London from a series of English newspapers (*General Evening Post*, *Lloyd's Evening Post*, *Lloyd's Lists*, *London Chronicle*, and *Middlesex Journal*).

Finally I use data on weather conditions from the observatory of Zwanenburg, a town close to Amsterdam (10 kms from the city centre). There are two or three observations a day on the wind direction and other weather variables. This data comes from *KNMI*.

2 Model

Under the null hypothesis all information, including private signals, is immediately incorporated into prices (see footnote 2 on page 2). Under the alternative hypothesis insiders trade strategically and private information is only slowly revealed. In this section I develop a simple Kyle (1985) model to analyze what the theoretical predictions are if insiders indeed trade strategically. In the next section I test these predictions empirically.

²³A future contract could have 4 possible expiration dates: February 15, May 15, August 15, or November 15. Prices reported were for the future contract ending at the nearest date.

²⁴There are some exceptions, if a boat arrived in Hellevoetsluys very early in the morning, it sometimes happened that the information from London was already available in Amsterdam on the same day. I use the publication dates of English news in the *Amsterdamsche Courant* and *Rotterdamsche Courant* to identify these cases.

The model features the trade of a single risky asset in two different markets, London (L) and Amsterdam (A). All relevant information originates in London. I fully abstract from public information (in the empirical setting it will be simple to reintroduce this) and focus on private signals. The full model consists of an infinite number of episodes, indexed with k . Each individual episode k is represented in figure 3. At the beginning of episode k nature determines the true value of the asset v_k , where v_k is a random walk, i.e. $v_k = v_{k-1} + \varepsilon_k$ with $\varepsilon_k \sim N(0, \sigma_\varepsilon^2)$. ε_k is not known to the wider public but is privately observed by a single agent, the London insider, at the beginning of the episode. At the end of the episode, ε_k is publicly revealed in London and the next episode $k + 1$ begins.

[FIGURE 3 ABOUT HERE]

The model is focussed on developments in Amsterdam. I assume that right after the moment nature decides on ε_k , before any trade takes place in London, the London insider sends this signal to a trusted agent in Amsterdam. This information arrives in Amsterdam at the beginning of period $t = 1$. When, at the end of the episode, ε_k is revealed, this information is also sent to Amsterdam immediately. Depending on the weather conditions, this news either arrives in Amsterdam relatively quickly after just one round of trade (referred to as $t = 1$), or it is delayed and only arrives after an additional round of trade ($t = 2$).²⁵ The probability of news arriving right after $t = 1$ is $1 - \pi_k$. The probability of it arriving after $t = 2$ is π_k . π_k is allowed to differ across episodes. The boat that leaves London after the episode has come to an end, fully reveals private signal ε_k . This means that after the arrival of the boat the informed agent loses his informational advantage. If news travels fast, there is therefore only one period to trade on private signal ε_k , whereas if news is delayed there are two periods.²⁶ Apart from π_k episodes are ex ante identical and I therefore drop subscript k .

Developments in Amsterdam are modeled as a two period Kyle (1985) model. There is a single risk neutral agent in Amsterdam who privately observes ε before any trade takes place. The only change with respect to Kyle (1985) is the introduction of uncertainty about whether the informed agent will have a second period to trade in or not (compare Back and Baruch 2004 and Caldenteu and Stacchetti 2010). Period $t = 2$ only occurs with probability π .²⁷

In addition to the informed agent, there is a competitive risk neutral market maker and every trading period features noise trading. Specifically, in both periods liquidity traders will submit an exogenous trading demand u_t with $u_t \sim N(0, \sigma_{u_t}^2)$. u_1 and u_2 are independent of each other and ε . I allow the variance of the liquidity shock to differ between periods.

The informed agent submits a market order x_t to the market maker. The market maker also

²⁵Period $t = 1$ of the model corresponds to periods in Amsterdam right after the arrival of a boat. Period $t = 2$ of the model corresponds to subsequent periods without news.

²⁶Note that the probability that a boat arrives after $t = 1$ or $t = 2$ also depends on the speed of the previous boat (the first solid line in figure 3 pointing to $t = 1$).

²⁷It is possible to add more periods to the model. However, this adds little to the model's key intuition. In addition, there are too few periods without news available in Amsterdam that are sufficiently long to estimate a multi-period model in the data.

receives the liquidity shock u_t as a market order. The market maker cannot discriminate between x_t and u_t and only observes $y_t = x_t + u_t$. The market maker is competitive and risk neutral and sets the price at which it processes orders at $p_1 = v_0 + E[v|y_1]$ and $p_2 = v_0 + E[v|y_1, y_2]$ (v_0 is the ex ante expected price, v_{k-1}). There are no constraints on the position the market maker can take.²⁸

As is standard in the literature I restrict the attention to linear equilibria.

Proposition 1 *A unique linear equilibrium exists and has the following form:*

$$x_1 = \beta_1 \varepsilon \tag{1}$$

$$x_2 = \beta_2 [\varepsilon - (p_1 - v_0)] \tag{2}$$

$$p_1 = v_0 + \lambda_1 (x_1 + u_1) \tag{3}$$

$$p_2 = p_1 + \lambda_2 (x_2 + u_2) \tag{4}$$

with

$$\beta_1 = \frac{1}{2\lambda_1} \frac{\lambda_2 - \frac{1}{2}\pi\lambda_1}{\lambda_2 - \frac{1}{4}\pi\lambda_1} \tag{5}$$

$$\beta_2 = \frac{1}{2\lambda_2} = \sqrt{\frac{\sigma_{u_2}^2}{(1 - \lambda_1\beta_1)\sigma_\varepsilon^2}} \tag{6}$$

$$\lambda_1 = \frac{\beta_1\sigma_\varepsilon^2}{\beta_1^2\sigma_\varepsilon^2 + \sigma_{u_1}^2} \tag{7}$$

$$\lambda_2 = \frac{\beta_2(1 - \lambda_1\beta_1)\sigma_\varepsilon^2}{\beta_2^2(1 - \lambda_1\beta_1)\sigma_\varepsilon^2 + \sigma_{u_2}^2} = \sqrt{\frac{(1 - \beta_1\lambda_1)\sigma_\varepsilon^2}{4\sigma_{u_2}^2}} \tag{8}$$

Proof. see appendix A. ■

This equilibrium is very similar to the one in Kyle (1985). The first key result is summarized by the following two corollaries (compare Chowdry and Nanda 1991).

Corollary 2 $cov(p_1 - v_0, \varepsilon) > 0$

Proof. see appendix A. ■

Corollary 3 $cov(p_2 - p_1, \varepsilon) > 0$

Proof. see appendix A. ■

²⁸Note that this specific market microstructure differs from the historical setting. An alternative setup would be a stylized two period limit order market (with risk averse agents) in the vein of Kyle (1989). However such a model adds significant complications. Numerical analysis indicates that results are qualitatively unchanged. For simplicity, I therefore stick to the Kyle (1985) framework.

These corollaries state that price changes in Amsterdam in both periods $t = 1$ and $t = 2$ should be correlated with the private signal ε . The monopolistic behavior of the insider leads to a slow revelation of the private signal. The insider takes his own price impact (Kyle's λ) into account and this constrains his behavior. In other words, the equilibrium of $t = 1$ is not fully revealing and asymmetric information is persistent. As a result there is additional price discovery going on in $t = 2$.

Price changes in London after the departure of a boat should also be correlated with ε . The private signal will be publicly announced in London by the time the next boat is set to depart for Amsterdam. Before that happens the London insider trades on his private information and ε will be (largely) revealed before the actual public announcement.²⁹ As a result, price changes in Amsterdam after the arrival of a boat should be positively correlated with the returns in London after the departure of that boat. Note that these price changes in London remain unreported in Amsterdam until the sailing of the next boat. This co-movement of Amsterdam and London prices is not driven by the transmission of public information. In other words, price changes in Amsterdam should be correlated with contemporaneous, but as of yet unreported, returns in London. This should both be observable for Amsterdam news returns (period $t = 1$ of the model) and Amsterdam non-news returns (period $t = 2$ of the model).

In line with Kyle's results, there is no auto-correlation between price changes in Amsterdam in $t = 1$ and $t = 2$. The intuition comes from market efficiency. If there is positive auto-correlation, the price change in $t = 1$ would not fully incorporate all relevant information. The market maker would be able to predict the price change in $t = 2$ based on the price change in $t = 1$. This would be inconsistent with risk neutral and competitive behavior.

The second key result is summarized by the following corollary

Corollary 4 $\frac{\delta cov(p_1 - v_0, \varepsilon)}{\delta \pi} < 0$

Proof. see appendix A. ■

This corollary states that the co-movement between the price change in $t = 1$ and ε should be decreasing in π . The intuition for this result follows from the trade-off the informed agent faces between profits from trading in $t = 1$ and $t = 2$. If the insider only has one period available to trade on his private signal, he would balance the price impact of his trade (the more he trades, the bigger the price impact) with the volume of trades he can get executed. One specific price impact-volume combination maximizes profits. If he gets a second period to trade, this optimal price impact-volume combination would change. The insider would now prefer to trade less in period $t = 1$ so that he reveals less of his private information and he can obtain more profits in period $t = 2$. In the

²⁹This process can be described by a 'standard' single or multi-period Kyle model. I assume that ε is publicly revealed by the time the next boat departs for Amsterdam. As a result the price discovery process in London has no impact on informed profits in Amsterdam - this simplifies the analysis.

model period $t = 2$ only occurs with probability π . If π is high the insider would like to trade less aggressively in $t = 1$ to save informational advantage for $t = 2$. However if π is small, the optimal strategy would be to trade more aggressively in $t = 1$. As a result, co-movement with ε (and thus with London) will be stronger.

What drives the variation of π ? First of all, π is determined by the speed of the boat that publicly reveals ε (see figure 3). Second, it is also driven by the speed of the boat that initially transmitted ε as a private signal. If this boat arrived relatively late to begin with, the (unconditional) probability of having two periods of trade is small.

There are a number of assumptions that merit further discussion. First of all, I assume that all private information is publicly revealed in London by the time the next boat sets its sails for Amsterdam. This is a simplification to make sure (1) there is only one private signal at a given point in time and (2) actions of the London insider have no impact on insider profits in Amsterdam. It is of course possible that the private signal is not revealed until later. This should not change the model's predictions. London prices should still reveal private information (albeit imperfectly) and the two markets should remain to move in the same direction. In addition, the Amsterdam insider would still trade more or less aggressively depending on when he expects news from London to arrive in Amsterdam. Even though this news is not fully revealing, it will still divulge part of the private signal.

Secondly, I assume that the London insider does not trade on a certain private signal before it is sent to Amsterdam. It is easy to allow for informed trading before the departure of a boat. As long as the private signal is not perfectly revealed the moment a boat departs for Amsterdam, this should not change the model's predictions.³⁰

Finally, I assume that the insider is a monopolist. There may in fact be multiple insiders. The model's predictions should be the same as long as the different private signals are sufficiently heterogeneous (Foster and Vishwanathan 1996). If multiple agents observe the same private signal, competition between these agents would quickly undo the model's predictions (Holden and Subrahmanyam 1994).

3 Empirical evidence

3.1 Introduction

In this section I present empirical evidence for the model's predictions. All estimates that are reported in the main text are for East India Company (EIC) stock only. Appendix B reports estimates for Bank of England (BoE) stock and the 3% Annuities.

To begin with, the model predicts that returns in Amsterdam should foreshadow developments

³⁰Under monopolistic behavior this is not likely to happen. The London insider will not trade too aggressively before the boat departs for Amsterdam. This would reduce overall insider profits.

in London, as the same underlying private signal gets incorporated into prices in both cities. This should be true for both $t = 1$, the news period right after the arrival of a boat, and $t = 2$, the subsequent non-news period. Secondly, the model predicts that this co-movement with London in period $t = 1$ should be stronger if the next boat is expected to arrive relatively quickly.

To guide the empirical discussion, figure 4 applies the model setup (figure 3) to the empirical setting. There are two relevant boat crossings. The first crossing transmits private signal ε_k (and publicly reveals previous private signal ε_{k-1}). The second crossing reveals private signal ε_k (and brings in new private signal ε_{k+1}). With probability $1 - \pi_k$ the second boat will arrive relatively quickly after the first, with probability π_k it will arrive relatively late.

[FIGURE 4 ABOUT HERE]

I define returns in London and Amsterdam as follows. R_k^L is the return in London that takes place after the departure of the first boat. I refer to this as the London post-departure return. To be clear, information about this return is not publicly transmitted by the first boat. I calculate this return over different periods (2, 3, 4, and 5 days). I use the 3 day return as a baseline - results for 2, 4 and 5 day London returns are presented in appendix B. $R_{k,t=1}^A$ is the return in Amsterdam that takes place immediately after the arrival of this boat. I refer to this as the Amsterdam news return. In terms of the model, this is the return that takes place over period $t = 1$. Since I only have 3 prices a week available for Amsterdam (Monday, Wednesday and Friday), this return is calculated over 2 or 3 day periods, depending on the day of the week.³¹ Finally, $R_{k,t=2}^A$, if it occurs, is the return in Amsterdam that takes place over the subsequent period without news, the so-called Amsterdam no-news return. This corresponds to period $t = 2$ of the model. Returns are calculated as log returns in percentages.

3.2 Re-introducing public information

So far I have completely abstracted from public information. It is clear that the first and the second boat in figure 4 do not only carry information pertaining to the private signals, they also carry public information. This includes that part of the previous private signal that is not yet incorporated into Amsterdam prices. It also captures all other public developments in London. Public news is captured by ΔN_{k-1} which measures the difference between the price in London right before the departure of a boat (p_{k-1}^L) and the Amsterdam price right before the arrival of a boat (p_{k-1}^A). This should capture all private and public information not yet incorporated into Amsterdam prices. Independent of whether private information is important and the Kyle model is valid or not, the Amsterdam news returns ($R_{k,t=1}^A$) should respond to this public news shock.

I control for these public information shocks throughout the empirical analysis. The most impor-

³¹The returns that are calculated for Friday to Monday (3 instead of 2 days) are not scaled down. The reason for not doing so that is that trading was restricted in the weekend. Jewish traders would not trade on the Sabbath, while Christians would not on Sundays (Spooner 1983). As an approximation, I therefore treat the 3 day weekend return the same as the 2 day week returns.

tant reason for doing so is that there might be momentum in the London return series that drives the co-movement between London and Amsterdam. Suppose that there are trading costs or limits to arbitrage that lead to a delayed response to public information and consequently to momentum (Hong and Stein 2001). In that case, we would expect Amsterdam and London prices to move in the same direction even if there is no private information at all. Controlling for public news shocks should control for this. In one of the robustness checks I also consider the case where trading costs or limits to arbitrage lead to zero returns instead of momentum.

3.3 Baseline results

How well do the corollaries established by the model hold up? I first consider the impact of insider trading and the resulting co-movement of London and Amsterdam returns. Corollaries 2 and 3 state that returns in Amsterdam, both right after the arrival of news from London and during subsequent periods without any news, should foreshadow developments in London that will take place right after the departure of the news to Amsterdam. Or in terms of figure 4 both $R_{k,t=1}^A$ (the Amsterdam news return) and $R_{k,t=2}^A$ (the Amsterdam no-news return) should be correlated with R_k^L (the London post-departure return). The regressions of interest are

$$R_{k,t=1}^A = \beta_0 + \beta_1 R_k^L + \beta_2 \Delta N_{k-1} + u_t \quad (9)$$

$$R_{k,t=2}^A = \beta_0 + \beta_1 R_k^L + \beta_2 \Delta N_{k-1} + v_t \quad (10)$$

Figures 5 and figures 6 present these correlations graphically for EIC stock. Figure 5 shows that $R_{k,t=1}^A$ is positively correlated with London return R_k^L (calculated over the three days after the boat departure). Likewise, figure 6 shows that $R_{k,t=2}^A$ is also correlated with R_k^L .³² This implies that price changes in Amsterdam predict the contemporaneous (but as of yet unreported) return in London – news of which will only arrive in the future.

[FIGURES 5 AND 6 ABOUT HERE]

In table 1 I estimate these correlations in a formal econometric framework, again for EIC stock. I correlate R_k^L (again calculated over the three days after the departure of the boat) with $R_{k,t=1}^A$ and $R_{k,t=2}^A$. I condition on the public news shock ΔN_{k-1} . This measures the public news shock arriving in Amsterdam with the packet boat. Inclusion of this variable should correct for any momentum that might be present in the return series.

[TABLE 1 ABOUT HERE]

Results in table 1 suggest that the co-movement with London post-departure returns is especially strong for returns taking place right after the arrival of a boat from England. However, the difference between the two coefficients is not significant at standard confidence levels. In addition, the reverse

³²Note that returns are corrected for the public news shock ΔN_{k-1} . To facilitate interpretation, the x-axes in figures (5) and (6) are truncated. See figures (15) and (16) in Appendix B for the raw plots. The regression estimates use all available data.

seems to be true for BoE stock and the 3% Annuities (see tables 7 and 8 in Appendix B), although differences here are not statistically significant either. It can be shown that the model has no strict predictions about this dimension; it all depends on the variance of the noise trading shocks in the two periods.

I perform two simple robustness tests. First of all, I calculate the London post-departure return R_k^L for different periods (2, 4 and 5 days after a boat departure). Results for EIC stock are presented in table 9 in Appendix B. Varying the period over which to calculate R_k^L does not matter substantially for the estimates. In addition, I exclude the first day of the London post-departure return to make sure that the positive co-movement between R_k^L and $R_{k,t=1}^A$ is not just driven by the fact that the boat carried slightly more recent public information than indicated by the data. Table 10 in Appendix B indicates that this is not the case.

3.4 Different expectations about the next boat arrival

Corollary 4 states that the co-movement with London in period $t = 1$ should be stronger if the next boat is expected to arrive relatively quickly. How well does this corollary hold up? This would provide direct evidence in favor of the strategic behavior of insiders.³³

Before going into the empirical tests let me first discuss how I estimate the crucial parameter π_k . This π_k is the probability that a second period of trade ($t = 2$) occurs within episode k (see figure 4). I allow π_k to differ over time. To generate an empirical equivalent of π_k I estimate $E[A]$. This is the expected number of days until the arrival of the next boat. This expectation is constructed right after the arrival of the previous boat. The longer $E[A]$, the more trading opportunities the informed agent has. I calculate this expectation in two ways.

To begin with, I take the median travelling time of a packet boat (4 days, including the day of departure) and add the number of days until the next departure from London is scheduled (or I subtract the number of days since departure if this has already happened). For example, suppose that a boat has just arrived and that the next boat is scheduled to depart from London one day later. I add one day to the median travelling time to arrive at the expected time until the next boat arrival. If this boat had left London one day earlier I subtract one day. This yields $E[A|simple]$. Two boats a week were set to sail between England and Holland, so the unconditionally expected number of days until the next arrival is 3.5. I differentiate between $E[A|simple] > 3.5$ and $E[A|simple] < 3.5$.

As an alternative, I allow the expected travelling time to vary depending on weather conditions and the time of the year. This yields $E[A|extended]$. Specifically I estimate a duration model with a flexible Gamma distribution that predicts travelling times. I condition on multiple factors. Most importantly, sailing boats had trouble crossing the North Sea from England to Holland when the

³³The model also holds predictions for co-movement in period $t = 2$. Specifically, we would expect to observe less co-movement if the insider initially expected the next boat to arrive right after $t = 1$. Unfortunately there are too few observations available to test this. It happens only infrequently that period $t = 2$ takes place unexpectedly.

wind was blowing from an eastern direction. When the sailing direction gets too close to the wind, sails cannot be adjusted anymore and a sailing boat enters the so-called no-go zone; see figure 7. If the boat's direction lies within the no-go zone, it will have to tack or, in other words, it will constantly have to change direction. This slows down sailing and leads to a longer effective distance. I have data available on wind directions from the observatory of Zwanenburg (close to Amsterdam) for 2 or 3 observations a day. For every observation I determine whether a sailing boat sailing east would face a no-go zone. For modern sailing boats this no-go zone lies around 30 to 50 degrees from the wind-direction. I assume that the 18th century packet boats had a no-go zone of 55 degrees around the prevailing wind direction. For every day I calculate what fraction of available observations featured a no-go zone, $\{0, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, 1\}$. I include this information as dummy variables. Other variables I condition on are a dummy for low temperatures to capture the possible presence of ice.³⁴ In addition, I include month dummies to condition on the season.³⁵

[FIGURE 7 ABOUT HERE]

The correlation between the actual number of days until the next boat arrival and $E[A|simple]$ and $E[A|extended]$ is 0.36 and 0.52 respectively. Figures 8 and 9 present these results graphically. Both figures plot Kaplan-Meier estimates of the fraction of boats that have not yet arrived t days after the arrival of the previous boat. I differentiate between $E[A]$ higher or lower than 3.5 days. If the estimation was perfect, we would expect all observations with $E[A] > 3.5$ to be still "at risk" at 3.5 days. All observations with $E[A] < 3.5$ should have "failed" by 3.5 days. Obviously this is not the case. However, both the simple and the extended procedures come up with reasonable estimates of $E[A]$, where the extended model does slightly better.

[FIGURES 8 AND 9 ABOUT HERE]

Now that we have defined $E[A|simple]$ and $E[A|extended]$, we can move to testing corollary 4. In table 2 I estimate whether Amsterdam news returns ($R_{k,t=1}^A$) on EIC stock co-move more strongly with London post departure returns (R_k^L) if the next news shipment is expected to arrive within less than 3.5 days. The regression of interest is

$$R_{k,t=1}^A = \beta_0 + \beta_1 R_k^L + \beta_2 R_k^L \times (E[A] < 3.5) + \beta_3 \Delta N_{k-1} + \beta_4 (E[A] < 3.5) + u_t$$

The interaction effect captures corollary 4. I use both $E[A|simple]$ and $E[A|extended]$. In both cases the interaction effect is positive and statistically significant at the 1 or 5% level. Economically the effects are also significant. Co-movement roughly doubles if the next boat is expected to arrive within 3.5 days. (Figures 17 to 20 in appendix B present these estimates graphically.) As expected

³⁴Hellevoetsluys was situated in the mouth of several rivers. Ice floating downstream could make it hard to reach the harbor. A dummy for temperatures below 3 degrees Celcius has the best fit in the duration model. This is the case because inland tempartures were lower than the ones measured at the Zwanenburg observatory which was relatively close to the sea.

³⁵In unreported results I find that the volatility of stock prices in London does not significantly differ across months. This suggests that these month dummies do not capture a different underlying price process.

the interaction effect is largest when $E[A|extended]$ is used. Tables 11 and 12 in Appendix B present these estimates for BoE stock and the 3% Annuities. Results are qualitatively similar. Table 2 calculates London returns over 3 days after the departure of a boat. In table 13 in Appendix B I redo these estimates, using different periods over which to calculate the London post-departure return (2, 4, 5 days). Results are very similar. All in all, these estimates are supportive of the model's prediction that co-movement should be stronger if the insider expects to have less time to benefit from his private signal.

[TABLE 2 ABOUT HERE]

3.5 Feedback effects and price discovery in London

So far the analysis has focused on Amsterdam. However, the presence of private information also implies a number of testable predictions for price discovery in London. Specifically, there may be feedback effects where Amsterdam price changes have an impact on London. If there is a noisy price discovery process taking place in both markets, and noise is not perfectly correlated, then the price signals from the two markets together are more informative than the price signals individually (Chowdry and Nanda 1991; Boulatov et al. 2011). In other words, London could update its beliefs based on Amsterdam price discovery. In the theoretical model of section 3 I assume that the private signal ε in London is revealed relatively quickly. This is an approximating assumption to facilitate the analysis of the price discovery process in Amsterdam. It is possible that private information in London was longer lived. If that is true, then the revelation of private information in Amsterdam could facilitate price discovery in London.

Figure 10 illustrates this situation. At London time t^* a boat with private signal ε sets sail to Amsterdam. There this information is received at Amsterdam time t . a days later a boat departs again for London, carrying information about the Amsterdam price at time $t+a$. The concrete price signal this boat carries can be expressed as $p_{t+a}^A - p_{t^*}^L$. This information reaches London at time t^*+l . Between t^* and t^*+l the London market observed a London signal that led to the price change $p_{t^*+l}^A - p_{t^*}^L$. At time t^*+l the London market maker updates his beliefs based on the two available signals; $p_{t^*+l}^A - p_{t^*}^L$ and $p_{t+a}^A - p_{t^*}^L$. This will lead to a new London price $p_{t^*+l+1}^L$.

[FIGURE 10 ABOUT HERE]

In what follows I analyze this situation statistically. Appendix C develops the statistical intuition into a full-fledged theoretical model. I show that under a number of reasonable assumptions predictions are equivalent.

Suppose that the Amsterdam and London signals have the following form

$$\begin{aligned}\theta^A &= \varepsilon + \zeta \\ \theta^L &= \varepsilon + \phi\end{aligned}$$

where everything is normally distributed. I assume that the Amsterdam and London signal are

positively correlated; $0 < cov(\theta^A, \theta^L) < var(\theta^i)$ for $i = A, L$. Assuming away any additional innovations in the price process (later on I will return to this assumption), price changes in Amsterdam and London can be written as

$$\begin{aligned} p_{t^*+l}^L - p_{t^*}^L &= \rho^L \theta^L \\ p_{t+a}^A - p_{t^*}^L &= \rho^A \theta^A \\ p_{t^*+l+2}^L - p_{t^*}^L &= \rho^{A|L} \theta^A + \rho^{L|A} \theta^L \end{aligned}$$

where

$$\rho^i = \frac{cov(\varepsilon, \theta^i)}{var(\theta^i)} = \frac{\sigma_\varepsilon^2}{var(\theta^i)} \quad (11)$$

$$\rho^{i|j} = \rho^i - \rho^{j|i} \frac{cov(\theta^i, \theta^j)}{var(\theta^i)}, \text{ for } i, j = A, L \quad (12)$$

Again, assuming away any additional shocks, this implies that

$$\begin{aligned} \theta^L &= \frac{p_{t^*+l}^L - p_{t^*}^L}{\rho^L} \\ \theta^A &= \frac{p_{t+a}^A - p_{t^*}^L}{\rho^A} \end{aligned}$$

and that

$$\begin{aligned} p_{t^*+l+1}^L - p_{t^*+l}^L &= \rho^{A|L} \theta^A - (\rho^L - \rho^{L|A}) \theta^L \\ &= \frac{\rho^{A|L}}{\rho^A} (p_{t+a}^A - p_{t^*}^L) - \frac{(\rho^L - \rho^{L|A})}{\rho^L} (p_{t^*+l}^L - p_{t^*}^L) \end{aligned} \quad (13)$$

Under the assumption that $0 < cov(\theta^A, \theta^L) < var(\theta^i)$, equation 13 has three testable predictions.

Prediction 1: *The price change in London ($p_{t^*+l+2}^L - p_{t^*+l}^L$) should respond positively to price changes in Amsterdam: $\frac{\rho^{A|L}}{\rho^A} > 0$. In addition, ($p_{t^*+l+1}^L - p_{t^*+l}^L$) should respond negatively to the previous London price change: $(\rho^L - \rho^{L|A}) > 0$*

Proof. see appendix A. ■

The intuition for this results follows from simple bivariate statistics. If the Amsterdam and London signal are both correlated with ε , and are not perfectly correlated with each other, then the London market maker can learn from the Amsterdam signal. At the same time, the London market maker will put less weight on the London signal in $t^* + l + 1$ (after the arrival of the Amsterdam boat) than he did in $t^* + l$ (before the arrival of the boat). He does so because, by assumption, price changes in Amsterdam and London are positively correlated and putting more weight on one signal automatically implies putting less weight on the other. The second part of this result is not trivial. If price changes in Amsterdam would simply reflect some fundamental news unrelated to private

information, there should be no reason why London price changes would partially revert after the arrival of a boat from Amsterdam.

Prediction 2: *In the empirical estimation $\frac{\rho^{A|L}}{\rho^A}$ will be larger when the response of $(p_{t^*+l+1}^L - p_{t^*+l}^L)$ on $(p_{t+a}^A - p_{t^*}^L)$ is made conditional on the London price change $(p_{t^*+l}^L - p_{t^*}^L)$.*

Proof. follows from omitted variable bias. ■

The intuition for this result again follows from $(\rho^L - \rho^{L|A}) > 0$ and the fact that, by assumption, price changes in London and Amsterdam are positively correlated. If one of these variables is omitted in the regression analysis this will lead to a downward bias of $\frac{\rho^{A|L}}{\rho^A}$. This is not a trivial result either. If London prices would simply respond to news about fundamentals from Amsterdam that is unrelated to the private information, the size of $\frac{\rho^{A|L}}{\rho^A}$ should not depend on the inclusion of $(p_{t^*+l}^L - p_{t^*}^L)$.

Prediction 3: *Finally, $\frac{\rho^{A|L}}{\rho^A}$ should be decreasing in the precision of signal θ^L (3a) and increasing in the precision of signal θ^A (keeping $cov(\theta^A, \theta^L)$ constant) (3b).*

Proof. See appendix A. ■

It is easy to see why $\frac{\rho^{A|L}}{\rho^A}$ drops when θ^L becomes more informative. ρ^A is unaffected, but $\rho^{A|L}$ will fall as more weight will be put on θ^L . The second part of the result is less intuitive. When θ^A becomes less precise then both ρ^A and $\rho^{A|L}$ should fall - overall there is less to learn from θ^A . The result shows that the drop in $\rho^{A|L}$ dominates. This, again, follows from $cov(\theta^A, \theta^L) > 0$. If there is only one signal available and this signal becomes less informative, then the weight on that signal (ρ^A) falls proportionally. When there are two correlated signals and one of these signals becomes less informative, then it is optimal to decrease the weight on that signal ($\rho^{A|L}$) more than proportionally.

Figure 10 illustrates how the precision of the Amsterdam and London signals vary over time and how they can be measured. First of all, when the time in Amsterdam between the arrival of news and the sailing of the next boat to London (a) takes relatively long, then Amsterdam prices are probably more informative. More trade has taken place and most likely more of the private signal has been revealed. Longer periods a occur when news from London happens to arrive right after a boat has just set sail for England. At that point it will take at least 3 or 4 days for the next boat to sail out. If weather conditions are such that boats cannot set sail for England, period a will take even longer. In the same vein, if the time in London between the initial departure of the boat and the eventual arrival of news from Amsterdam (l) is long, then London prices are more informative. In this case London prices probably reveal a large part of the private information. Longer periods l occur when a is long (creating a identification problem, see discussion below), or when sailing times on the North Sea (in either direction) happened to be long.

To test these predictions I estimate the following regression, including interaction effects between

Amsterdam returns and a and l to pick up the effect of the two different signals' precision.

$$\begin{aligned}
p_{t^*+l+1}^L - p_{t^*+l}^L &= \beta_0 + \beta_1 (p_{t+a}^A - p_{t^*}^L) + \beta_2 (p_{t^*+l}^L - p_{t^*}^L) \\
&+ \beta_3 (p_{t+a}^A - p_{t^*}^L) \times a + \beta_4 (p_{t+a}^A - p_{t^*}^L) \times l \\
&+ \gamma_1 a + \gamma_2 l + u_t
\end{aligned} \tag{14}$$

Predictions 1 and 3 imply that $\beta_1 > 0$, $\beta_2 < 0$, $\beta_3 > 0$ and $\beta_4 < 0$. Before turning to the regression results, let me first reiterate that I make an important assumption to arrive at (13) and (14), namely that no additional shocks affect the price process. This is obviously not true. Price change $(p_{t+a}^A - p_{t^*}^L)$ may include additional noise. In addition, $(p_{t^*+l}^L - p_{t^*}^L)$ may also include new private and public information shocks. In other words, I only observe θ^L and θ^A with an error. As a result there will be attenuation bias in estimating the regression coefficients from (14). However, predictions 1-3 should still remain valid. Coefficients (and differences between coefficients) should simply be smaller.

[TABLE 3 ABOUT HERE]

Table 3 presents the regression results from estimating 14 step by step.³⁶ Results are for EIC stock. In the first column I estimate β_1 ($\frac{\rho^{A|L}}{\rho^A}$) individually. There seems to be some response in London to developments in Amsterdam. However, consistent with the impulse responses in figure 14 in Appendix B, β_1 is small. Unconditionally, price changes in Amsterdam have little impact on prices in London.

Consistent with prediction 2, this β_1 increases with about 50% to 0.1 when the previous return in London is included in the estimation (a 1% increase in the Amsterdam price leads to a 0.1% increase in London prices). This difference is statistically significant at the 1% level. Consistent with prediction 1 β_2 ($-\frac{(\rho^L - \rho^{L|A})}{\rho^L}$) is negative and significant. These results suggest that the London reaction to news from Amsterdam is not just simply related to some other fundamental shocks originating in Holland but is really the result of the feedback effect of private information.

Columns 3 and 4 include the interaction effects with a and l individually. Both a and l are introduced as deviations from the median. This means that coefficient β_1 is estimated at median values for a (3 days) and l (11 days). The interaction effects measure the impact on β_1 when moving away from this median. Neither interaction effect, economically nor statistically, is significant on its own. This is not unexpected. a and l are positively correlated (correlation of 0.73). The two hypothesized interaction effects actually have opposite expected coefficients ($\beta_3 > 0$ and $\beta_4 < 0$). Introduced individually, the coefficients will be biased downwards. In column 5 they are introduced

³⁶The observations included in these regressions are constructed as follows (see figure 10 for reference). For every news shipment from Amsterdam that arrived in London, I check when this news left Amsterdam. I then check what the most recent date of the English news was in Amsterdam the moment this boat departed. I calculate a , the number of days between the arrival of this news and the departure of the present boat. I then finally check on what date this English news had been sent from London. That way I can calculate l . Prices $p_{t^*}^L$, p_{t+a}^A , $p_{t^*+l}^L$ and $p_{t^*+l+1}^L$ are easy to recover.

jointly. In this specification they do have statistically significant coefficients. The signs on the coefficients are as predicted. A longer l leads to a smaller response of the London price to Amsterdam price changes. A longer a leads to a larger response.

The economic impact of the interaction effects is considerable. The 90th percentile of the distribution of $[a - \text{median}(a)]$ is at 2 days. This means that moving from the median to the 90th percentile increases the response coefficient from 0.111 to 0.245. The 10th percentile of the distribution of $[l - \text{median}(l)]$ is at -3 days. This means that moving from the median to the 10th percentile increases the response coefficient from 0.111 to 0.264.

Tables 14 and 15 in Appendix B present these estimates for BoE stock and the 3% Annuities. Results are qualitatively similar, although the economic effects are smaller than in the EIC case.

To summarize, the empirical evidence is consistent with Amsterdam price discovery affecting London prices. This effect does not seem to be driven by public news shocks originating in Amsterdam that are unrelated to private information. London's response to Amsterdam increases substantially after conditioning on past London price changes. This is inconsistent with independent public news shocks coming from Amsterdam. In addition, the effect is significantly larger when there was more time to trade in Amsterdam or when the overall time the private signal needed to "bounce off" was relatively short. Again, this is consistent with a feedback effect of private information, but inconsistent with the arrival of public news from Amsterdam.

4 Robustness checks

The baseline findings of section 3.3 can be driven by other factors than private information. In this sub-section I discuss three alternative explanations: (1) the slipping through of public news through alternative channels, (2) the slow incorporation of public information into prices and (3) correlated liquidity shocks. Note that none of these alternative explanations can explain why co-movement between Amsterdam was stronger when the next boat was expected to arrive relatively early. Nor can they explain the feedback effect of Amsterdam prices on London. However they might affect the baseline estimates.

4.1 Slipping through of news

The historical record does not suggest that alternative channels through which English news could reach Amsterdam played an important role. Hope & Co for example, the most important Anglo-Dutch bank of the period (and involved in insider trading), fully relied on the packet boat service. Hope did not hire its own private boat to get information from England. Nevertheless, it is impossible to rule out that others did. The key question is whether the (possible) slipping through of news can fully explain the co-movement results or whether there remains a role for private information.

To answer this question I use the fact that sailing boats relied on the weather to get across the

North Sea. I restrict the sample to periods where, after the arrival of a packet boat, wind conditions suddenly turned so that future packet boats were significantly delayed. I assume that during these periods it was equally impossible for other boats to get across. If co-movement was purely driven by news arriving through other channels, prices should not move in the same direction in London and Amsterdam during these episodes. If there is evidence for co-movement, this would not necessarily disprove that channels other than the packet boats were used. These could still be relevant during good weather conditions. However it would indicate that private information is needed to fully explain the patterns in the data.

The potential slipping through of information is most relevant for Amsterdam non-news returns ($R_{k,t=2}^A$).³⁷ I construct three different definitions of bad weather. First of all I distinguish between no-news periods that were purely the result of the sailing schedule and those that were not and must have been the result of bad weather (definition A). I use the scheduled departure dates of the packet boats and the median sailing time to reconstruct by when a boat should have arrived if weather conditions were normal. If by that time no boat had arrived, I count this as a bad weather episode. Bad weather sample B is based on wind directions. Returns in Amsterdam are measured over 2 or 3 day periods. For every return I determine what the average daily wind direction was during this 2 or 3 day period. If the average wind condition was east (from 0 to 180 degrees) on every single day, I include the observation in bad weather sample B. I do something similar for bad weather sample C. Here I look at the no-go zones (see figure 7). For every day of a 2 or 3 day period, I check how many wind observations within that day (out of a total of 2 or 3) featured a no-go zone. If for every day at least 2 of these daily wind observations featured a no-go zone, I include this return in bad weather sample C.

[FIGURES 11 TO 13 ABOUT HERE]

Figures 11 to 13 plot EIC non-news returns $R_{k,t=2}^A$ against London post-departure returns R_k^L for the three restricted bad weather samples. The figures show that during periods when it was difficult to get news across the North Sea, there was still a positive correlation between returns in London and Amsterdam. The relevant regression is given by:

$$R_{k,t=2}^A = \beta_0 + \beta_1 R_k^L + \beta_2 R_k^L \times badweather + \beta_3 \Delta N_{k-1} + \beta_4 badweather$$

where I use all three different bad weather definitions. The interaction term between R_k^L and the bad weather dummy captures whether there was a different degree of co-movement during episodes of adverse weather conditions.

[TABLE 4 ABOUT HERE]

Table 4 presents the corresponding results. The table shows that the degree of co-movement was equally strong in periods of bad weather - the interaction term is economically small, statistically

³⁷ Amsterdam news returns took place right after the arrival of an official boat and is unlikely that a non-official boat both arrived around the same time and contained more recent information. For this to happen, the non-official boats would have had to be faster than the official packet boats.

insignificant and has a positive (instead of negative) sign for two of the three bad weather definitions. Tables 16 and 17 in Appendix B present the same estimates for BoE stock and the 3% Annuities and results are almost identical. In table 18 in Appendix B I redo these estimations for EIC stock, calculating London returns over 2, 4 or 5 days (instead of 3 days) after the departure of a boat, using bad weather definition A. Again, results remain virtually the same.

To summarize, these results indicate that even under adverse weather conditions, when official boats were unable to sail, there was still co-movement between London and Amsterdam prices. This does not conclusively rule out that public news could have reached Amsterdam through alternative channels during better weather conditions. However it does indicate that private information is needed to fully explain the patterns in the data.

4.2 Slow incorporation of public news

It is possible that co-movement is driven by the slow incorporation of public news into prices in London. One possible reason for this could be trading costs or limits to arbitrage. If it is costly to trade on new public information, prices might not be fully updated. This slow incorporation of public news could be accompanied by momentum (see Hong and Stein 2001 and references therein). All estimates are corrected for this. Results indicate that although return continuation is present, it does not drive the co-movement results.

Alternatively, it is possible that trading costs or limits to arbitrage in London are so significant that (occasionally) prices do not adjust at all to reflect new information (Lesmond et al. 1999; Bekaert et al. 2007). This would lead to zero returns. Information would be incorporated into London prices at a later point. This could potentially cause the co-movement patterns between Amsterdam and London that I document in the baseline regressions. I check whether co-movement between Amsterdam and London is stronger after zero pre-departure returns (R_{k-1}^L) in London. I run the following regressions for Amsterdam news and non-news returns:

$$\begin{aligned} R_{k,t=1}^A &= \beta_0 + \beta_1 R_k^L + \beta_2 R_k^L \times [R_{k-1}^L = 0] + \beta_3 \Delta N_{k-1} + \beta_4 [R_{k-1}^L = 0] + u_t \\ R_{k,t=2}^A &= \beta_0 + \beta_1 R_k^L + \beta_2 R_k^L \times [R_{k-1}^L = 0] + \beta_3 \Delta N_{k-1} + \beta_4 [R_{k-1}^L = 0] + v_t \end{aligned}$$

The interaction term between R_k^L and $[R_{k-1}^L = 0]$ captures whether there is a different degree of co-movement between $R_{k,t=1,2}^A$ and R_k^L if the London pre-departure return is zero.

[TABLE 5 ABOUT HERE]

Table 5 presents the regression estimates for EIC stock. Results indicates that for Amsterdam news returns there is a positive interaction effect. There is more co-movement if previous returns in London were zero. However this does not drive the co-movement results (compare table 5 with table 1. The size of the coefficient on R_k^L is hardly affected by the inclusion of $R_k^L \times [R_{k-1}^L = 0]$). There is no impact on Amsterdam no-news returns. Tables 19 and 20 in appendix B present similar

estimates for BoE stock and the 3% Annuities. For these securities there is no evidence for stronger co-movement after zero returns.

4.3 Permanent price changes or liquidity shocks?

The co-movement between Amsterdam and London could be driven by correlated liquidity shocks rather than by private information. Suppose a London agent is hit by a large liquidity shock. He may decide to split his orders between Amsterdam and London to minimize overall price impact. As he moves down (potentially) downward sloping demand curves in both markets, prices move in the same direction. How can we differentiate between this explanation and the slow revelation of private information? It is natural to assume that liquidity shocks have a transitory impact on prices (Grossman and Miller 1988). In other words, liquidity shocks should lead to return reversals. If this drives the co-movement between Amsterdam and London, then we should observe return reversals across markets. I.e. positive (negative) price changes in London should predict subsequent negative (positive) returns in Amsterdam.

To test this, I run a series of regressions where I predict Amsterdam price movements based on past returns in London. Specifically, the past London return is defined as the return before the sailing of a boat (the London pre-departure return R_{k-1}^L , see figure 4). The price change in Amsterdam is calculated between the arrival of that boat ($p_{k,t=1}^A$, see figure 4) and the Amsterdam price T days into the future (p_{k+T}^A). T varies between 2/3 days and 4 weeks. Note that $p_{k,t=1}^A$ should incorporate all public information that is contained in R_{k-1}^L . The regressions therefore do not pick up the response to public news.

[TABLE 6 ABOUT HERE]

I present these estimates in table 6. The estimates show no significant predictive power. If anything, these estimates indicate the presence of return continuation or momentum (for which I correct in all estimates) rather than reversals. In other words, the evidence suggests that transitory price movements in London do not travel across the North Sea. The co-movement between Amsterdam and London reflects permanent price changes rather than transitory shocks. Tables 21 and 22 in Appendix B present similar estimates for BoE stock and the 3% Annuities.

5 Conclusion

This paper studies the effect of privately informed trading on security prices, using a natural experiment from history. The results strongly support the classical Kyle (1985) model of strategic insider trading. The evidence is based on the market for British securities in Amsterdam during the 18th century, when British securities were traded both there and in London. Anecdotal evidence suggests that London insiders traded in both markets to benefit from privileged information. To do so, they sent letters to their agents in Amsterdam who would then trade on their behalf. Letters were sent

via mail packet boats, which carried both public information and private letters. These boats only sailed twice a week, and in adverse weather they could not sail at all. As a result, Amsterdam was frequently cut off from London news. I exploit these periods of exogenous market segmentation to identify the impact of private information and the strategic behavior of insiders.

To guide the empirical discussion I use a two period model based on Kyle (1985). In the model an informed agent trades slowly over time. Price changes in both periods reflect part of his signal. The rate at which private information is incorporated into prices depends on when the agent expects the private signal to be publicly revealed.

Empirical results are consistent with the model's predictions. Price movements in Amsterdam between the arrival of boats were correlated with the contemporaneous (but as of yet unreported) returns in London. This is consistent with the presence of a private signal that is slowly incorporated into prices in both markets. The initial co-movement of Amsterdam and London prices was stronger when the next boat was expected to arrive early. This is consistent with strategic behavior on the part of the insider. The next boat would carry public news that would reveal (part of) the private signal. If the insider had less time available to benefit from his informational advantage he would trade less aggressively early on. This accelerated price discovery.

The importance of private information is underlined by the response of London prices to price discovery in Amsterdam. Conditional on its own price discovery, the London market updated its beliefs based on price changes in Amsterdam. If due to weather conditions it took longer for a private signal to "bounce off" from Amsterdam, the London market responded less - by that time the private signal was already largely incorporated into London prices.

I provide evidence that the co-movement between Amsterdam and London reflected permanent price changes. This means that it is unlikely that the co-movement was the result of correlated liquidity trades or other transitory shocks. Robustness checks also suggest that the results are not driven by the slow incorporation of public information into prices due to trading costs or limits to arbitrage. Finally, the arrival of news through other channels than the official packet boats cannot account for the empirical results either.

18th century London and Amsterdam are in many ways the perfect testing ground for models of strategic insider trading. The key strength of this specific historical context is that information flows were less complex than today and can be perfectly reconstructed. Crucially, information arrived in a non-continuous way. This allows for a clean identification of private information. There are no confounding effects from other sources of information. In addition, the lengths of time over which insider information remained private in Amsterdam varied exogenously and this allows for a direct test of the strategic behavior of insiders.

Nevertheless, one might ask how general the results from this historical episode are. How crucial are the differences between then and now for the interpretation of the paper's findings? First of all, insider trading has become illegal since the 18th century and one might think that private information

has therefore become less relevant. This does not square well with the empirical evidence for private information in today's markets (see the literature overview on page 5). In addition, financial markets have tremendously increased in scale and depth since the 18th century. This would imply that today trading has become more anonymous and there might be even more opportunities for insider trading.

A second important difference lies in how quickly and frequently information is transmitted. Obviously, the transmission of information was a lot slower and more infrequent in the 18th century. The identification strategy of this paper depends on these characteristics. However, private information itself, the thing that we fundamentally care about, was not the result of primitive communication technology. Rather it was the result of London insiders having access to superior sources of information. This is similar to today where corporate insiders are likely to be better informed than the market as a whole.

Thirdly, the results of this paper are consistent with long-lived private information (I estimate that it took two weeks for a given signal to be incorporated into prices) and with insiders having monopoly power over their private signals.³⁸ This may be different today. Insiders could have become more competitive. Alternatively, due to more advanced technology, private information may have become much shorter lived. This would have fundamentally changed the role of private information in markets - they may have become much closer to strong form efficiency. However, recent research suggests that this is not the case. Even after decades of fast technological progress, markets have not become more informative. For example, earnings surprises are as big as they always were (Bai et al. 2012). This suggests that long lived private information, held by monopolistic agents, is as relevant today as it ever was.

³⁸Either private information was held by a single insider (Kyle 1985) or multiple insiders observed different signals (Foster and Viswanathan 1996).

References

HISTORICAL NEWSPAPERS:

Amsterdamsche Courant, City Archive Amsterdam
Opregte Haerlemsche Courant, Dutch Royal Library, The Hague
Rotterdamsche Courant, City Archive Rotterdam
General Evening Post, Burney Collection, British Library
Lloyd's Evening Post, Burney Collection, British Library
Lloyd's Lists
Middlesex Journal, Burney Collection, British Library
London Chronicle, Burney Collection, British Library

CITY ARCHIVES AMSTERDAM (SAA) AND ROTTERDAM (OSA):

Estate Uriel Cardozo, SAA 334, 643
Hope & Co, SAA 735, 78; 79; 115 and 1510 (letterbooks)

Stukken betreffende het postverkeer met Engeland, 1687, 1697-1710. [Documents regarding the postal connection with England, 1687, 1697-1710.], OSA 2599

LITERATURE

Admati, A.R. and P. Pfleiderer, 'A theory of intraday patterns: volume and price variability', in *The Review of Financial Studies* 1-1 (1988).

Back, K., 'Insider Trading in Continuous Time', in: *Review of Financial Studies* 5 (1992)

Back, K., and S. Baruch, 'Information in securities markets: Kyle meets Glosten and Milgrom', in: *Econometrica* 72-2 (2004).

Bai, J., T. Philippon and A. Savov, 'Have financial markets become more informative?', mimeo (2012)

Banerjee, S., R. Kaniel, and I. Kremer, 'Price drift as an outcome of differences in higher-order beliefs', in: *Review of Financial Studies* (2009)

Bekaert, G., C.R. Harvey and C. Lundblad, 'Liquidity and expected returns: lessons from emerging markets', in: *The Review of Financial Studies* 20-5 (2007)

Boulatov, A., T. Hendershott and D. Livdan, 'Informed trading and portfolio returns', *Review of Financial Studies* (2011, forthcoming)

Caldentey R. and E. Stacchetti, 'Insider trading with a random deadline', in: *Econometrica* 78-1 (2010).

Chau, M. and D. Vayanos (2008), 'Strong-form efficiency with monopolistic insiders' in: *The Review of Financial Studies* 21-5 (2008).

Chordia, T., R. Roll and A. Subrahmanyam, 'Order imbalance, liquidity and market returns', in: *Journal of Financial Economics* 65 (2002).

Chowdry, B. and V. Nanda (1991), 'Multimarket Trading and Market Liquidity', in *The Review of Financial Studies* 4-3, pp. 483-511

- Cohen, L., A. Frazzini and C. Malloy, 'The small world of investing: board connections and mutual fund returns', in: *Journal of Political Economy* (2008)
- Cohen, L., A. Frazzini and C. Malloy, 'Sell side school ties', in: *Journal of Finance* (2010)
- Cohen, L., C. Malloy and L. Pomorski, 'Decoding inside information', *Journal of Finance* (2011, forthcoming)
- Dahya, J., Z. Shi, and J. van Bommel (2010), 'An empirical investigation of the speed of information aggregation: a study of IPOs', in: *International Journal of Banking, Accounting and Finance* 2-1, pp. 47-79
- Dempster, G., J.M. Wells, and D.T. Wills, 'A common-features analysis of Amsterdam and London financial markets during the eighteenth century', *Economic Inquiry* (2000).
- Duarte, J., and L. Young, 'Why is PIN priced?', in: *Journal of Financial Economics* 91 (2009).
- Easley, D., N. Kiefer and M. O'Hara, 'One day in the life of a very common stock', in: *The Review of Financial Studies* 10-3 (1997)
- Evans, M.D.D. and R.K. Lyons, 'Order flow and exchange rate dynamics', in: *Journal of Political Economy* 110-1 (2002).
- Ferguson, N. 'The House of Rothschild. Money's Prophets 1798 - 1848', Viking Penguin (1998)
- Foster, F. Douglas, and S. Viswanathan, 'Strategic trading with asymmetrically informed traders and long-lived information', in: *Journal of Financial and Quantitative Analysis* 29-4 (1994).
- Foster, F. Douglas, and S. Viswanathan, 'Strategic trading when agents forecast the forecasts of others', in: *The Journal of Finance* 51-4 (1996).
- Glosten, L.R. and P.R. Milgrom, 'Bid, ask and transaction prices in a specialist market with heterogeneously informed traders', in: *The Journal of Financial Economics* 14 (1985).
- Grossman, S.J., and J. Stiglitz, 'On the impossibility of informationally efficient markets', in: *The American Economic Review* 70-3 (1980).
- Grossman, S.J., and M.H. Miller, 'Liquidity and market structure', in: *The Journal of Finance* 43-3 (1988).
- Grundy, B. D. and M. McNichols, 'Trade and the Revelation of Information through Prices and Direct Disclosure', in: *The Review of Financial Studies* 2-4 (1989)
- Hasbrouck, J., 'Measuring the information content of stock trades', in: *The Journal of Finance* 46-1 (1991).
- He, H. and J. Wang, 'Differential information and dynamic behavior of stock trading volume', in: *The Review of Financial Studies* 8-4 (1995)
- Hemmeon, J.C., *The History of the British Post office*, Cambridge: Harvard University Press (1912).
- Hendershott, T., D. Livdan and N. Schuerhoff, 'Are institutions informed about news?', mimeo (2012)
- Hoes, H.J., 'Voorgeschiedenis en ontstaan van het Financiële dagblad', in *Economisch-Historisch Jaarboek* 49 (1986).
- Hogesteeger, G., *Van lopende bode tot telematica*, Groningen: PTT Nederland (1989).
- Holden, C.W. and A. Subrahmanyam, 'Long-lived private information and imperfect competition', in:

The Journal of Finance 47 (1992).

Holden, C.W. and A. Subrahmanyam, 'Risk aversion, imperfect competition, and long-lived information', in: *Economics Letters* 44 (1994).

Hong, H., and J. Stein, 'Disagreement and stock market', in: *Journal of Economic Perspectives* 21-2 (2001)

Jansen, L., *De Koophandel van Amsterdam*, Nieuwe Uitgevers Maatschappij, Amsterdam (1946).

Jegadeesh, N., 'Evidence of predictable behavior of security returns', in: *The Journal of Finance* 45-3 (1990).

Jonker, J., *Merchants, bankers, middlemen*, Amsterdam: NEHA (1996).

Kelly, B. and A. Ljungqvist, 'Testing asymmetric-information asset pricing models', *Review of Financial Studies* (2011, forthcoming).

Koudijs, P., 'The boats that did not sail: news and asset price volatility in a natural experiment', working paper (2012)

Kyle, A.S., 'Continuous auctions and insider trading', in: *Econometrica* 53-6 (1985).

Kyle, A.S., 'Informed speculation with imperfect competition', in: *Review of Economic Studies* 56-3 (1989)

Lesmond, D.A., J.P. Ogden, and C.A. Trzcinka, 'A new estimate of transaction costs', in: *The Review of Financial Studies* 12-5 (1999)

Levi, W., *The Pigeon*. Sumter: Levi Publishing Co. (1977).

Llorente, Guillermo, Roni Michaely, Gideon Saar, and Jiang Wang (2002), 'Dynamic Volume-Return Relation of Individual Stocks', in: *The Review of Financial Studies* 15-4, pp. 1005-1047.

Madhavan A., M. Richardson and M. Roomans, 'Why Do Security Prices Change? A Transaction-Level Analysis of NYSE Stocks', in: *Review of Financial Studies* 10-4.

Neal, L., *The Rise of Financial Capitalism*, Cambridge: Cambridge UP (1990).

Pasquariello, P., and C. Vega, 'Informed and strategic order flow in the bond markets', *The Review of Financial Studies* (2007)

Polak, N.J., 'Het nooteeringsvraagstuk aan de Amsterdamsche effectenbeurs in de eertse helft der negentiende eeuw', in: *Economisch-Historisch Jaarboek* 10 (1924).

Rogers, Th., *Agriculture and prices in England*, Henry Frowde, Oxford, Vol. 7-2 (1902)

Romer, D., 'Rational Asset-Price Movements Without News', in: *The American Economic Review* 83-5 (1993).

Smith, M.F.J., *Tijd-affaires in effecten aan de Amsterdamsche beurs*, The Hague: Martinus Nijhoff (1919).

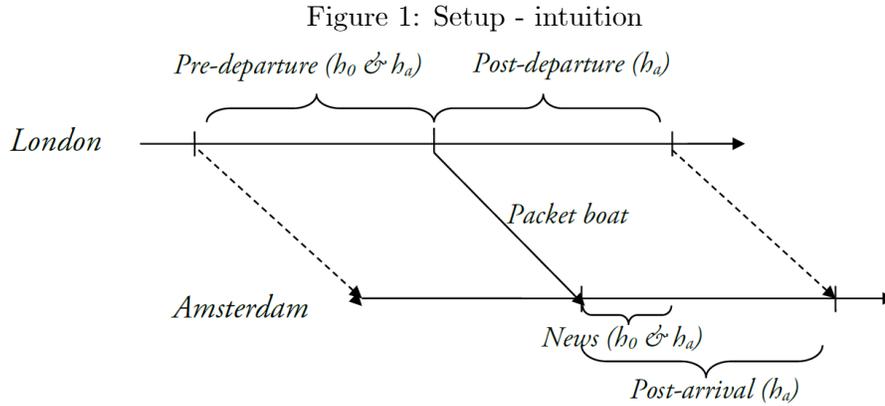
Spooner, F., *Risks at sea. Amsterdam insurance and maritime Europe*, Cambridge Cambridge UP (1983).

Stitt Dibden, W.G., *Four hundred years of Anglo-Dutch mail*, Postal History Society, The Hague (1965)

Subrahmanyam, Avanidhar (1991), 'Risk Aversion, Market Liquidity, and Price Efficiency', in: *The Review of Financial Studies* 4-3, pp. 417-441.

- Sutherland, L., *The East India Company in eighteenth-century politics*, Oxford: Oxford UP (1952).
- Ten Brink, E.A.B.J., *De geschiedenis van het postvervoer*, Bussum: Fibula (1969).
- Ten Brink, E.A.B.J., *De Rotterdamsche Telegraaf Maatschappij, 1854-1884 : een episode uit de begintijd van de elektrische telegrafie in Nederland*, The Hague (1957).
- Tetlock, P.C., ‘Does Public Financial News Resolve Asymmetric Information?’, in: *Review of Financial Studies* 23, pp. 3520-3557 (2010).
- Van Dillen, J.G., ‘Effectenkoersen aan de Amsterdamsche Beurs 1723-1794’, in: *Economisch Historisch Jaarboek* 17 (1931).
- Van Nierop, L., ‘Brieven van den Amsterdamschen makelaar Robert Hennebo aan den Haarlemschen zijdereder Simon Bevel over Engelschen Fondsen (1735-1736)’, in: *Economisch Historisch Jaarboek* 17 (1931).
- Vega, Clara, ‘Stock price reaction to public and private information’ in: *Journal of Financial Economics* 82, pp. 103–133 (2006)
- Vives, X., ‘Short-Term Investment and the Information Efficiency of the Market’ in: *The Review of Financial Studies*, 8-1, pp. 125-160 (1995)
- Wang, J., ‘A model of competitive stock trading volume’, in: *The Journal of Political Economy* 102-1(1994).
- Wilson, C., *Anglo-Dutch commerce & finance in the eighteenth century*, Cambridge, Cambridge University Press (1941,1966).

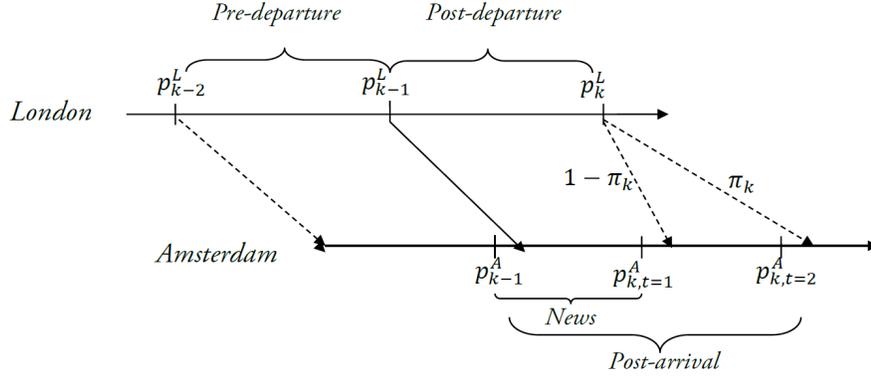
Figures



This diagram illustrates the identification strategy of this paper. The horizontal lines are time axes. The diagonal lines indicate the sailing of packet boats. Under the null hypothesis all information was public or, if there was private information as well, it was held by competitive agents and immediately incorporated into prices. This is captured by prices in London before the departure of a boat. Amsterdam prices should immediately reflect this information after the arrival of a boat (and the transmission of the news). Under this hypothesis prices in Amsterdam should only respond to price changes in London **before** the departure of a boat (h_0).

Under the alternative hypothesis additional co-movement takes place. Private information is held by strategic agents and is slowly incorporated into prices. The packet boat transmits this private information to insiders in Amsterdam. As insiders in London and Amsterdam trade on the same private information, prices in both markets should move in the same direction. This means that London post-departure and Amsterdam post-arrival returns should be positively correlated (h_a).

Figure 4: Setup - empirics



This diagram clarifies the timing of the empirical analysis and defines the different returns that are used. A given boat transmits information from London to Amsterdam. With probability π the next boat arrives quickly and there are limited opportunities to trade (period $t = 1$ of the theoretical model) and we only observe a “news return”. With probability $1 - \pi$ the next boat is delayed. Now there are additional opportunities to trade (period $t = 2$ in the theoretical model) and we also observe a “non-news return”. More precisely, returns are defined as follows:

$$\begin{aligned}
 \text{Amsterdam post-arrival news returns (} t = 1 \text{):} & \quad R_{k,t=1}^A = p_{k,t=1}^A - p_{k-1}^A \\
 \text{Amsterdam post-arrival non-news returns (} t = 2 \text{):} & \quad R_{k,t=2}^A = p_{k,t=2}^A - p_{k,t=1}^A \\
 \text{London post-departure returns:} & \quad R_k^L = p_k^L - p_{k-1}^L \\
 \text{London pre-departure returns:} & \quad R_{k-1}^L = p_{k-1}^L - p_{k-2}^L \\
 \text{Public news shocks (at the beginning of episode } k \text{):} & \quad \Delta N_{k-1} = p_{k-1}^L - p_{k-1}^A
 \end{aligned}$$

The paper’s analysis can be summarized by the following three statements:

$$\begin{aligned}
 \text{Under the null and alternative hypothesis:} & \quad cov(\Delta N_{k-1}, R_{k,t=1}^A) > 0 \\
 \text{Under the alternative hypothesis:} & \quad cov(R_k^L, R_{k,t=1}^A) > 0 \ \& \ cov(R_k^L, R_{k,t=2}^A) > 0 \\
 & \quad \& \frac{\partial cov(R_k^L, R_{k,t=1}^A)}{\partial \pi} < 0
 \end{aligned}$$

The regressions of interest are given by:

$$\begin{aligned}
 (1a) \quad R_{k,t=1}^A &= \beta_0 + \beta_1 R_k^L + \beta_2 \Delta N_{k-1} + u_t && \text{(benchmark news)} \\
 (1b) \quad R_{k,t=2}^A &= \beta_0 + \beta_1 R_k^L + \beta_2 \Delta N_{k-1} + v_t && \text{(benchmark non-news)} \\
 (2) \quad R_{k,t=1}^A &= \beta_0 + \beta_1 R_k^L + \beta_2 R_k^L \times (\pi_k < \bar{\pi}) + \beta_3 \Delta N_{k-1} + \beta_4 (\pi_k < \bar{\pi}) + w_t && \text{(small vs large } \pi \text{)}
 \end{aligned}$$

Figure 5: Co-movement LND post-departure and AMS post-arrival news returns

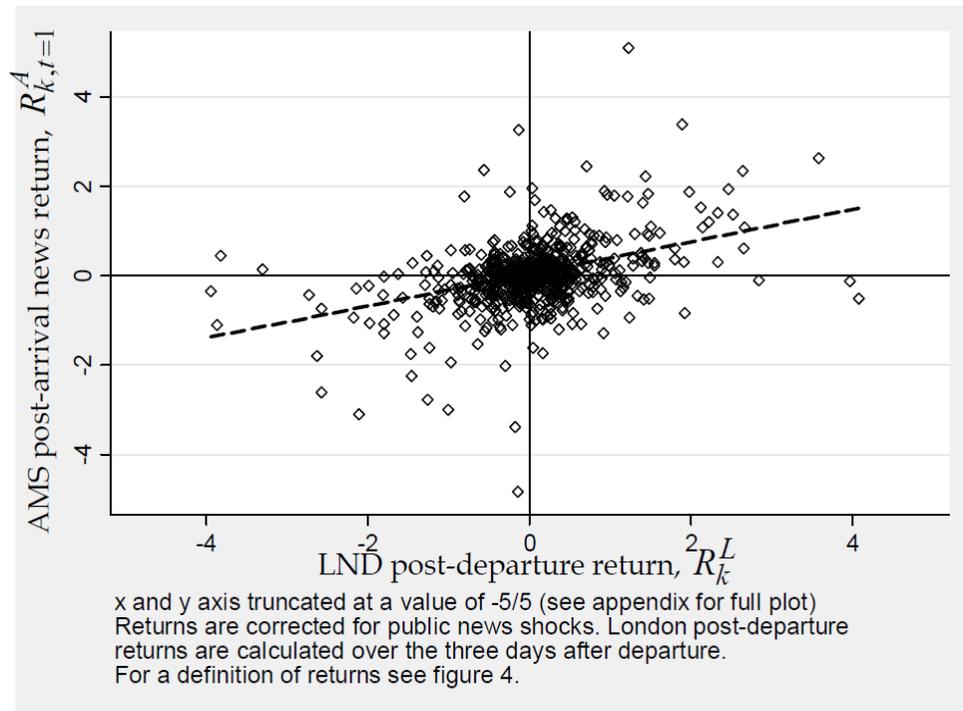


Figure 6: Co-movement LND post-departure and AMS post-arrival non-news returns

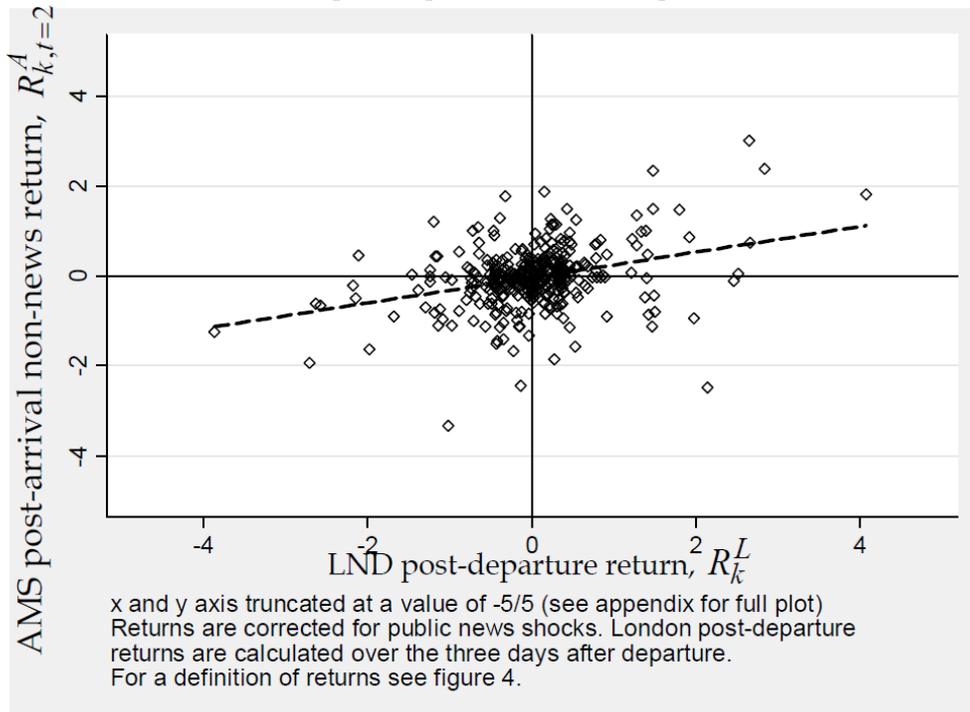


Figure 7: Points of sail (shaded area is the no-go zone)

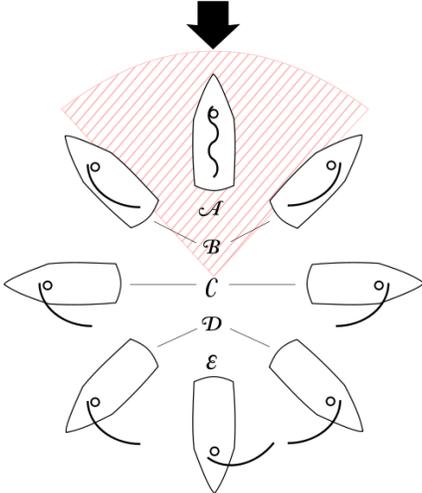


Figure 8: Kaplan-Meier estimates - arrival next boat (simple model)

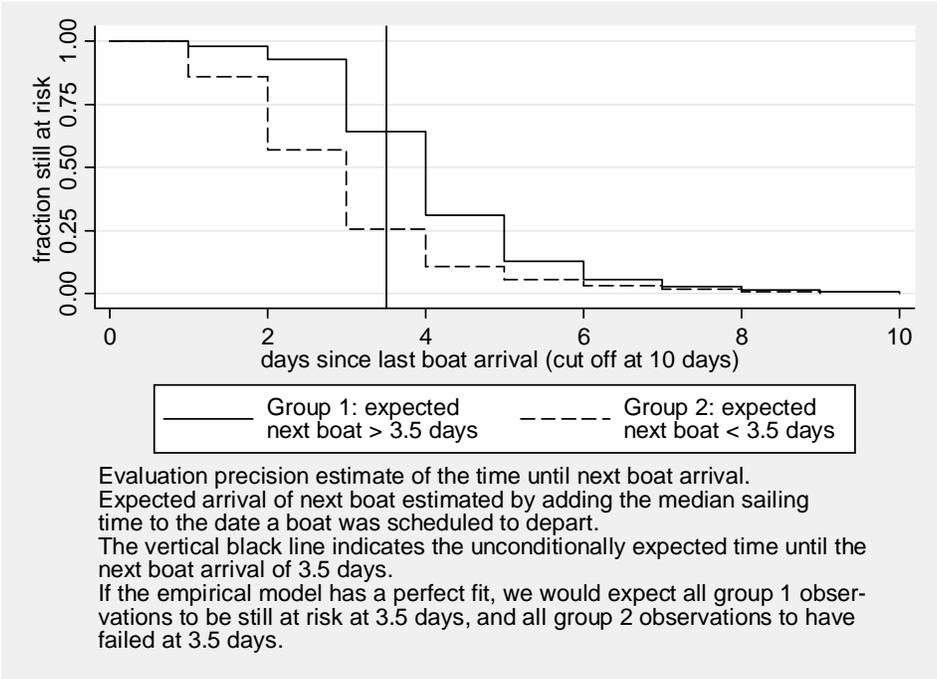


Figure 9: Kaplan-Meier estimates - arrival next boat (extended duration model)

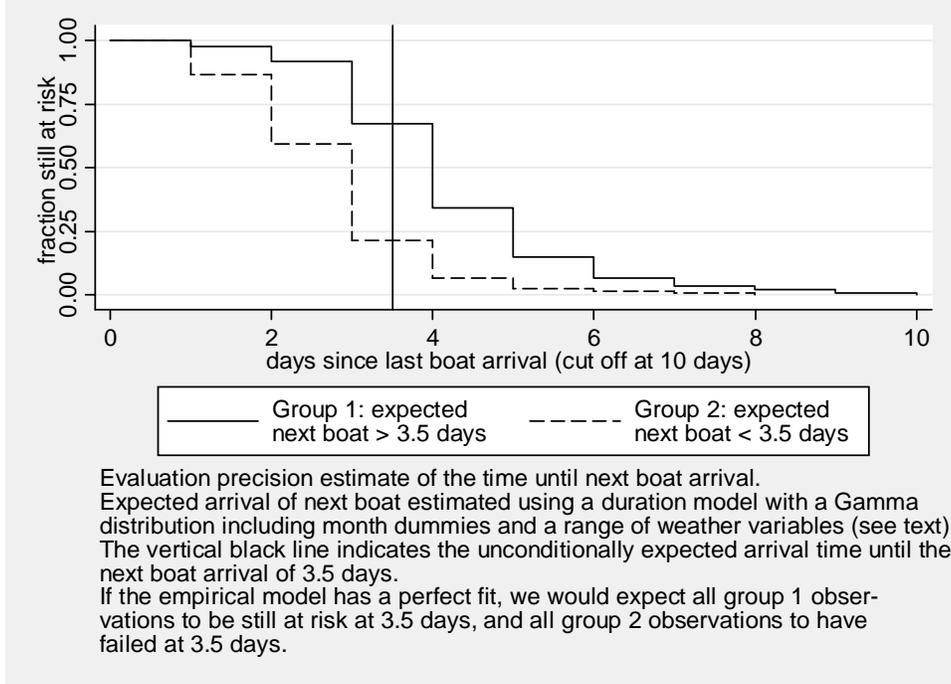
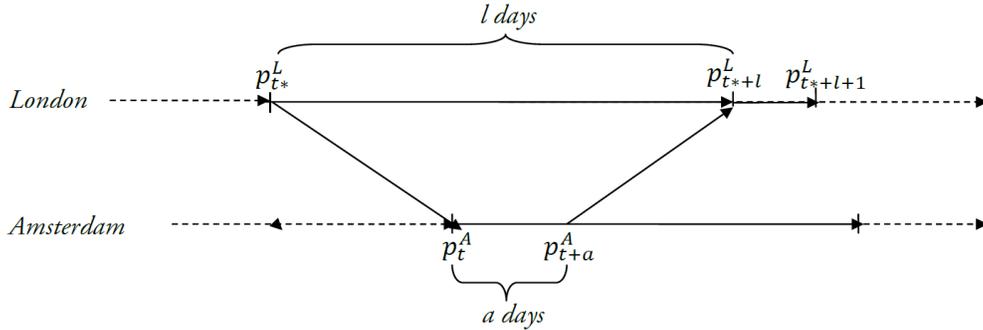


Figure 10: Setup - feedback effects



This diagram illustrates the feedback effect from Amsterdam back to London. A certain private signal leaves London at time t^* and arrives in Amsterdam at time t . There are a days between the arrival of this signal in Amsterdam and the departure of the next boat to London. Once this boat finally arrives in London l days have been passed since the signal was originally sent. Variation in a and l is determined by weather conditions at sea.

The relevant returns are given by:

$$\begin{aligned} \text{Amsterdam news returns:} & \quad p_{t+a}^A - p_{t^*}^L \\ \text{London pre-news returns:} & \quad p_{t^*+l}^L - p_{t^*}^L \\ \text{London news returns:} & \quad p_{t^*+l+1}^L - p_{t^*+l}^L \end{aligned}$$

The relevant regression is given by:

$$p_{t^*+l+1}^L - p_{t^*+l}^L = \beta_0 + \beta_1(p_{t+a}^A - p_{t^*}^L) + \beta_2(p_{t^*+l}^L - p_{t^*}^L) + \beta_3(p_{t+a}^A - p_{t^*}^L) \times a + \beta_4(p_{t+a}^A - p_{t^*}^L) \times l + \gamma_1 a + \gamma_2 l + u_t$$

Figure 11: Co-movement LND post-departure and AMS post-arrival non-news returns - bad weather (A)

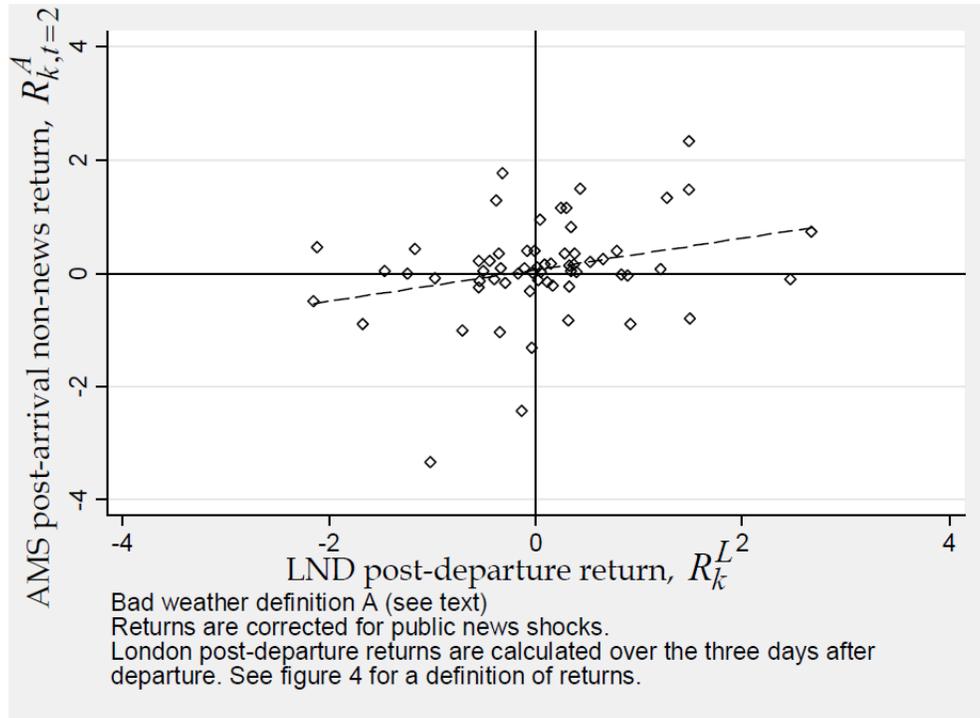


Figure 12: Co-movement LND post-departure and AMS post-arrival non-news returns - bad weather (B)

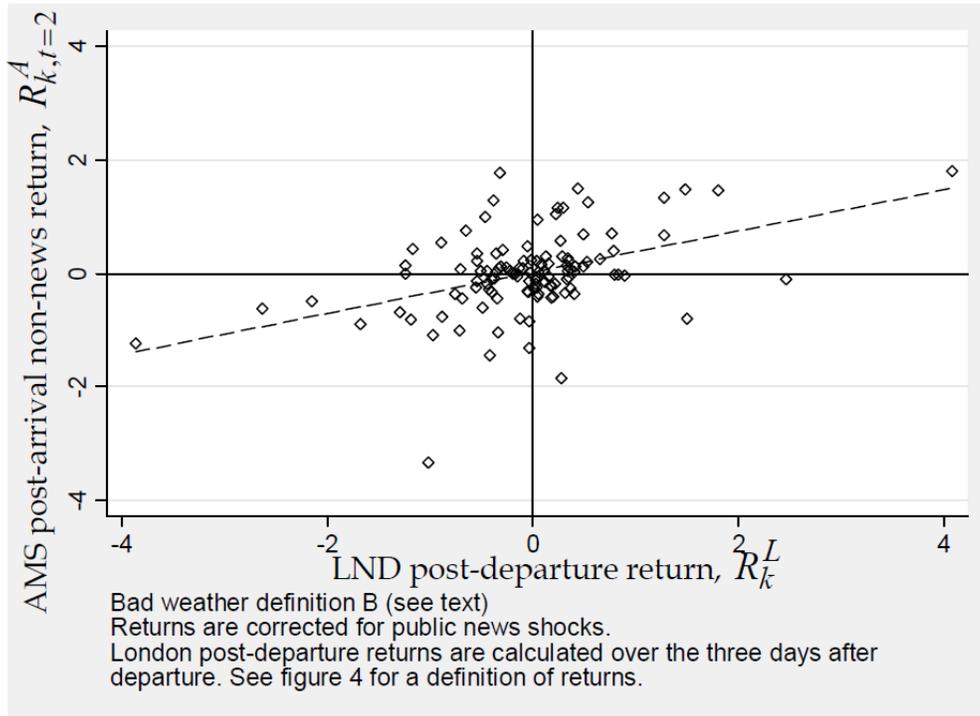
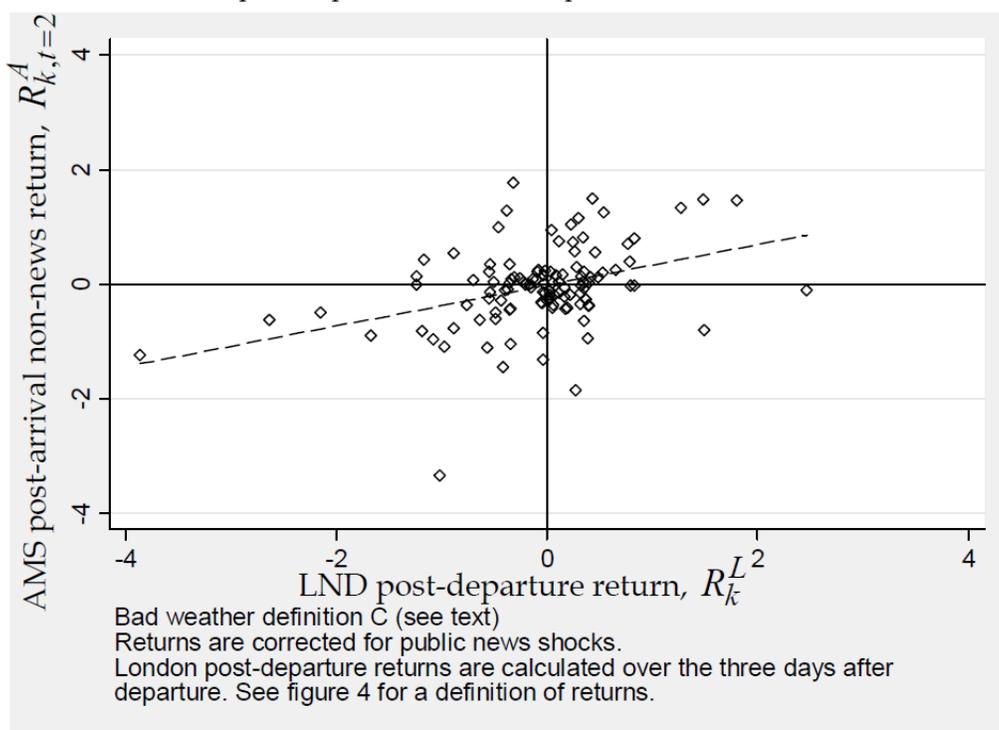


Figure 13: Co-movement LND post-departure and AMS post-arrival non-news returns - bad weather (C)



Tables

Table 1: Co-movement of returns - EIC

	(1)	(2)
	AMS post-arrival news return, $R_{k,t=1}^A$	AMS post-arrival non-news return, $R_{k,t=2}^A$
LND post-departure return, R_k^L	0.338 (0.048)***	0.220 (0.062)***
Public news shock ΔN_{k-1}	0.441 (0.040)***	0.077 (0.038)**
Constant	0.005 (0.026)	-0.023 (0.030)
Obs	666	355
Adj. R2	0.36	0.10
Chi^2 test (p-value)	2.27 (0.132)	

Estimates of co-movement between London post-departure and Amsterdam post-arrival returns. Estimates are adjusted for the possible continuation (momentum) of public news by including ΔN_{k-1} .

See figure 4 for exact definitions of returns. London post-departure returns are calculated over the three days after a boat departure. A Chi^2 test is performed on the equality of the R_k^L coefficients in columns (1) and (2).

***, ** denotes statistical significance at the 1, 5% level. Robust, bootstrapped standard errors (1000 replications) are reported in parentheses.

Table 2: Co-movement EIC, different expectations next boat

	AMS post-arrival news return, $R_{k,t=1}^A$	
London post-departure return, R_k^L	0.256 (0.060)***	0.251 (0.058)***
$R_k^L \times E[A simple] < 3.5$	0.223 (0.101)**	
$E[A simple] < 3.5$	-0.054 (0.060)	
$R_k^L \times E[A extended] < 3.5$		0.277 (0.091)***
$E[A extended] < 3.5$		-0.031 (0.055)
Public news shock, ΔN_{k-1}	0.479 (0.037)***	0.473 (0.039)***
Constant	0.021 (0.036)	0.007 (0.038)
Obs	627	627
Adj. R2	0.39	0.39

Estimates of co-movement between London post-departure and Amsterdam post-arrival returns. $E[A]$ stands for the expected number of days until the next boat arrival. $E[A|simple]$ is calculated by adding the median sailing time to the departure date of the next boat. For $E[A|extended]$ the median sailing time is replaced by a conditionally expected sailing time which is estimated in a duration model using a Gamma distribution, including a number of weather variables and month dummies (see main text).

See figure 4 for exact definitions of returns. London post-departure returns are calculated over three days after a boat departure. The observation of November 20, 1772 is dropped from the regression analysis to make sure that this outlier does not drive the positive interaction effect (see figures 17 and 19 in appendix B). Inclusion of this datapoint leads to slightly higher estimates of the interaction effect.

***, **, * denotes statistical significance at the 1, 5, 10% level. Robust, bootstrapped standard errors (1000 replications) are reported in parentheses.

Table 3: Feedback effects - EIC

	London returns between $t^* + l + 1$ and $t^* + l$ ($p_{t^*+l+1}^L - p_{t^*+l}^L$)				
	(1)	(2)	(3)	(4)	(5)
Amsterdam news returns ($p_{t+a}^A - p_{t^*}^L$)	0.066 (0.034)*	0.100 (0.039)***	0.100 (0.038)***	0.107 (0.039)***	0.111 (0.039)***
London pre-news returns ($p_{t^*+l}^L - p_{t^*}^L$)		-0.073 (0.041)*	-0.073 (0.041)*	-0.072 (0.041)*	-0.066 (0.041)*
Amsterdam period a			-0.027 (0.023)		-0.046 (0.035)
$a \times (p_{t+a}^A - p_{t^*}^L)$			0.011 (0.019)		0.067 (0.031)**
London period l				-0.004 (0.013)	0.022 (0.020)
$l \times (p_{t+a}^A - p_{t^*}^L)$				-0.013 (0.011)	-0.051 (0.019)***
<i>constant</i>	0.054 (0.037)	0.049 (0.036)	0.042 (0.037)	0.050 (0.036)	0.043 (0.038)
χ^2 test ($p_{t+a}^A - p_{t^*}^L$): (1) = (2) (p-value)		2.89 (0.089)*			
N	696	695	695	695	695
Adj. R^2	0.01	0.02	0.02	0.02	0.03

This table provides estimates of the feedback effect of Amsterdam returns in London prices. See figure 10 for a definition of the timing and the returns. a measures the number of days between the arrival of a signal in Amsterdam and the departure of the next boat to London. l measures the number of days it takes for the private signal to “bounce off” from Amsterdam. Variation in a and l is driven by weather conditions. All estimates, including the benchmark coefficient on $(p_{t+a}^A - p_{t^*}^L)$ are at median values of l (11 days) and a (3 days).

*, **, *** indicate significance at the 1, 5, and 10% level.

Table 4: Co-movement EIC, bad weather

	AMS post-arrival no-news return, $R_{k,t=2}^A$		
London post-departure return, R_k^L	0.258 (0.074)***	0.223 (0.089)***	0.266 (0.073)***
$R_k^L \times badweather(A)$	0.012 (0.148)		
$badweather(A)$	0.119 (0.111)		
$R_k^L \times badweather(B)$		0.115 (0.116)	
$badweather(B)$		0.067 (0.072)	
$R_k^L \times badweather(C)$			-0.020 (0.154)
$badweather(C)$			0.130 (0.094)
Public news shock, ΔN_{k-1}	0.120 -(0.067)***	0.126 (0.045)***	0.120 (0.043)***
Constant	-0.068 (0.034)	-0.067 (0.040)	-0.071 (0.037)
Obs - total	363	363	363
Obs - $badweather(A)$	65		
Obs - $badweather(B)$		117	
Obs - $badweather(C)$			67
Adj. R2	0.12	0.12	0.12

Estimates of co-movement between London post-departure and Amsterdam post-arrival non-news returns, conditional on good or bad weather conditions. See the text for a description of the three definitions (A, B, and C) of bad weather.

See figure 4 for exact definitions of returns. London post-departure returns are calculated over three days after a boat departure.

*** denotes statistical significance at the 1% level. Robust, bootstrapped standard errors (1000 replications) are reported in parentheses.

Table 5: Co-movement after zero returns in London, EIC

	Amsterdam post-arrival news return, $R_{k,t=1}^A$	Amsterdam post-arrival non-news return, $R_{k,t=2}^A$
London post-departure return, R_k^L	0.314 (0.055)***	0.218 (0.067)***
Zero London pre-departure return ($R_{k-1}^L = 0$)	-0.065 (0.064)	0.147 (0.074)**
× London post-departure return, R_k^L	0.183 (0.092)	0.034 (0.135)
Public news shock, ΔN_{k-1}	0.474 (0.039)***	0.050 (0.036)
Constant	0.012 (0.033)	-0.047 (0.034)
Obs - total	628	327
Obs - zero returns	108	58
Adj. R2	0.39	0.10

This table provides additional tests whether co-movement was driven by trading costs or limits to arbitrage. Results in preceding tables show that co-movement was not driven by return continuation. This table tests whether co-movement between Amsterdam and London was stronger if past London returns (pre-departure returns $R_{k-1}^L = 0$) had been zero. This proxies for situations where trading costs or limits to arbitrage led to a delay in the incorporation of public information into prices (Lesmond et al 1999).

See figure 4 for exact definitions of returns. London post-departure returns are calculated over three days after a boat departure.

***, * denotes statistical significance at the 1, 5 and 10% level. Robust, bootstrapped standard errors (1000 replications) are reported in parentheses.

Table 6: Permanent price changes? EIC

	Amsterdam return over period T after episode k ($p_{k+T}^A - p_{k,t=1}^A$)					
	<i>2/3 days</i>	<i>4/5 days</i>	<i>1 week</i>	<i>2 weeks</i>	<i>3 weeks</i>	<i>4 weeks</i>
London return (R_{k-1}^L)	0.043 (0.043)	0.028 (0.060)	0.077 (0.060)	0.077 (0.077)	0.087 (0.104)	0.195 (0.115)*
Constant	-0.007 (0.028)	0.011 (0.038)	0.068 (0.048)	0.131 (0.067)*	0.159 (0.087)	0.214 (0.102)
N	734	733	731	726	719	717
Adj. R^2	0.00	0.00	0.00	0.00	0.00	0.00

Estimates of regressions predicting future Amsterdam returns based on London returns.

The London return is defined as the London pre-departure return (R_{k-1}^L , see figure 4).

The Amsterdam return is calculated as the price change after the arrival of the boat that brings this information and the Amsterdam price T days into the future ($p_{k+T}^A - p_{k,t=1}^A$).

T varies between 2-3 days and 4 weeks.

Robust, bootstrapped (1000 reps) standard errors in parentheses.

* indicates statistical significance at the 10% level