Beyond Words: The Differential Impact of Fed Chairs' Facial Expressions on Financial Markets*

Hunter Ng
City University of New York
hunterboonhian.ng@baruch.cuny.edu

4th May 2025

Abstract

This paper explores how Federal Reserve Chairs' facial expressions during FOMC press conferences influence investor behavior and financial markets. Using facial recognition technology and deepfake simulations on press conference videos from April 2011 to December 2020, I quantify changes in nonverbal signals while controlling for verbal content. My findings reveal that nonverbal cues act as independent public signals that significantly affect market outcomes. Using deepfakes. I uniquely demonstrate that identical facial expressions elicit different market reactions and this depends on the Fed Chair's identity, tenure, and experience, indicating that investor interpretations are dynamically shaped by perceptions of the Chair. Moreover, the evolving market response over time aligns with the dual-processing, bounded memory model of information processing. Lastly, I find no evidence that Fed Chairs strategically change their facial expressions to influence markets, highlighting the unintentional yet impactful nature of nonverbal communication.

Keywords: Federal Reserve, Facial Expressions, FOMC Press Conferences, Nonverbal Communication, Deepfake

JEL Classifications: E44, E52, E58, G14, G41

*I am grateful to Lin Peng, Sophia Kazinnik, Svenja Dube, Monica Neamtiu for helpful comments. I also received valuable feedback from presentation participants at Baruch College, City University of New York.

1. Introduction

Nonverbal communication by the Fed Chair is an important investor information source which affects investor sentiments even when controlling for the verbal content of FOMC press conferences (Curti and Kazinnik, 2023). In today's zeitgeist of artificial intelligence (AI) and deepfakes, the question remains as to whether the Chair's facial expressions could be strategically controlled or simulated by deepfakes to produce better market outcomes. In this article, I attempt to fill this gap through the use of deepfakes and facial emotional recognition (FER) technology.

In today's modern era of central bank transparency, effective communication by the Federal Reserve is crucial for shaping market expectations and delivering information to markets (Woodford, 2001; Yellen, 2016). The introduction of post-FOMC press conferences in 2011 marked a significant development in how the Federal Reserve communicates, providing a platform for clarifying policy decisions and conveying the underlying motivations behind them. The primary purpose of FOMC press conferences is to communicate the Fed's policy intentions transparently and to reduce market uncertainty (Bernanke, 2013). Perceived transparency during FOMC press conferences is shown to be beneficial, as it aligns market expectations with the central bank's dual mandate (Geraats, 2009). However, beyond the carefully crafted verbal messages, studies show that the unintentional yet impactful nature of the Fed Chair's nonverbal cues also influence market perceptions.

While the verbal content of FOMC statements has a significant informational effect on market prices(Gómez-Cram and Grotteria, 2021), facial expressions offer a separate, fundamental source of information (Darwin, 1872). Advances in FER technology have expanded research into nonverbal communication cues, with studies by Curti and Kazinnik (2023) and Alexopoulos et al. (2024) documenting a correlation between market performance and the Fed Chair's facial expressions during events such as the FOMC press conferences and congressional testimonies. Given that facial expressions influence market reactions, I analyze how such nonverbal communication can be optimized.

Using a ten-year period of press conference videos from April 27, 2011, to December 16, 2020, I develop a novel methodological approach that combines FER technology (Kaur et al., 2022) with deepfake simulations to isolate and quantify nonverbal signals while controlling for verbal content. To do so, I use deepfake techniques to create counterfactual scenarios where identical facial expressions are overlaid across different Fed Chair identities, and examine how perceptions of nonverbal cues vary by Chair characteristics such as identity, tenure, and experience. Firstly, I find that a realistic deepfake of Fed Chair A making the same expressions as Fed Chair B registers different emotional readings using FER technology. This methodology advances existing research by demonstrating how deep learning tools can be used to disentangle the effects of identity from expression in high-stakes financial communication settings.

Secondly, I find that investors' interpretation of facial signals is dependent on the Fed Chair's tenure and experience, and that the frequency of negative expressions by Fed Chairs increases as their tenure lengthens. Investors also get more used to the Fed Chair's expressions over time and do not react as negatively to negative emotions. However, recent events such as congressional testimonies before the FOMC press conference increases the investors' reactions. This indicates a learning effect where investors interpret Fed Chair's emotions based on their past experience and recent events. This result supports the theory that investors take a bounded Bayesian approach when interpreting nonverbal cues (Wilson, 2014).

Lastly, I find no evidence that Fed Chairs strategically alter their facial expressions to deliberately influence market outcomes. This novel finding highlights the inherent, human aspect of Fed Chairs' participation in FOMC press conferences, suggesting that it is precisely this authentic and unintentional nature of nonverbal communication that allows investors to continue to extract meaningful information. Consequently, this helps explain the observed persistence of nonverbal cues' effects, even after controlling rigorously for verbal content.

This article is related to three strands of literature. Firstly, there is a growing economics empirical literature examining information signals of the Federal Reserve using sophisticated technologies and high-frequency data (Gomez-Cram and Grotteria, 2022; Curti and Kazinnik, 2023; Alexopoulous et al., 2024; Swanson and Jayawickrema, 2023; Gorodnichenko et al., 2023). These papers use novel ways of converting non-numerical information such as textual, nonverbal or other types of data in the Federal Reserve's arsenal to examine their effects on information transmission, which is revolutionizing the way we understand how investors process macroeconomic information.

Next, there is a stream of financial economics literature on how investors incorporate signals. Traditionally, central banks have used full-information rational expectations models to guide monetary policy in the wake of the Lucas's (1972) imperfect information model (Calvo, 1983; Coibion, Gorodnichenko and Kamdar, 2018). However, Gómez-Cram and Grotteria (2022) compare multiple theoretical frameworks and find that the difference-in-opinion model by Allen, Morris, and Shin (2006) best explains investors' interpretations of public signals from FOMC communications.

Lastly, this article relates to the use of deepfakes in interdisciplinary research. Westerlund (2019) provides a thorough review of how deepfakes can be used in society. The hyper-realistic videos are created using AI and can digitally recreate actual people, for example, the Fed Chair giving a speech. Renier et al., 2024 show that deepfakes can be used to recreate nonverbal behavioral studies while Emett et al., 2024 find that investors react to deepfake financial news using a realism heuristic, where they cannot properly differentiate between a super realistic deepfake and the real video.

This paper makes three contributions to the literature on financial communication, investor learning, and methodological innovation in economics. Firstly, it introduces a novel methodological framework that leverages deepfake technology combined with facial emotion recognition (FER) tools to counter-factually analyze the impact of nonverbal emotional signals. By creating highly realistic simulations in which identical facial expressions are overlaid onto different Fed Chair identities, this approach isolates the causal effects of expressions independent of verbal content or identity. This methodology offers researchers a new tool to disentangle communicative signals and demonstrates the potential of deep learning technologies for counterfactual economic analysis.

Second, the paper documents that Fed Chairs may not strategically manipulate their facial expressions to influence market outcomes. The findings reveal that nonverbal cues, while unintentional, nonetheless serve as impactful and independent signals affecting investor behavior.

This insight underscores the inherent human element in Federal Reserve communications and suggests that transparency goals may inadvertently leave room for emotional inference by market participants, a consideration increasingly important as the Federal Reserve seeks alternative communication levers (Bernanke, 2007).

Third, the paper documents that investors' reactions to nonverbal signals are dynamic and adaptive. While initial market responses to negative emotional expressions are pronounced, these effects diminish over time as investors gain experience with a particular Fed Chair's baseline emotional style. However, recent salient events, such as congressional testimonies, can reactivate heightened sensitivity to emotional cues. This pattern supports a Bayesian learning model of investor behavior (Wilson, 2014) and advances the understanding of how bounded memory and experience can be a starting point for studying how investors interpret nonverbal cues.

The remainder of this paper is structured as follows: Section 2 reviews relevant literature on central bank communication and nonverbal cues, and develops the hypotheses. Section 3 describes the data and methodology. Section 4 presents the empirical results, and Section 5 concludes with a summary of my key insights, caveats and suggestions for policy-making.

2 Hypothesis Development

2.1. FOMC Press Conference and Non-verbal Communication

The FOMC Committee holds eight meetings a year to discuss monetary policy actions. Since May 1999, the FOMC Committee started issuing post-FOMC meeting statements which specified target levels for the federal funds rate. After the federal funds rate hit 0% in December 2008, former Chair Ben Bernanke decided in 2011 to give press conferences after select meetings as an additional policy tool, and since 2019, every FOMC meeting has been followed by a press conference. These post-FOMC press conferences are the focus of this study, as they provide a rich source of both verbal and nonverbal communication from the Fed Chair.

The market responds significantly to post-FOMC press conferences. Lucca and Moench (2015) document a pre-FOMC meeting stock price drift, while Boguth et al. (2019) show that this price drift occurs only when a post-FOMC press conference is scheduled. Amengual and Xiu (2018) find that market volatility decreases after the announcement, indicating that information released during these conferences reduces uncertainty. The overwhelming evidence demonstrates that post-FOMC press conferences convey crucial information to the markets.

Moreover, the market response is not limited to before or after the press conferences; it occurs concurrently during the events. Gómez-Cram and Grotteria (2022) timestamp the words spoken during the press conferences and align them with high-frequency financial data, revealing a positive correlation between changes in the newly issued policy statement and stock returns. This shows that investors process information in real-time, reacting to both the content and delivery of the Fed Chair's communication.

Throughout FOMC press conferences, facial expressions serve as a critical channel for conveying emotional states, significantly shaping the reception of the Federal Reserve's messages. Psychological research underscores the primacy of nonverbal communication: Mehrabian's (1972) seminal "7-38-55 rule" suggests that merely 7% of a message is transmitted through words, whereas vocal elements and nonverbal cues such as facial expressions account for 38% and 55%, respectively. Recent advances in facial emotion recognition (FER) technologies have further enhanced the ability to systematically detect and classify emotional signals. For instance, Zhang et al. (2023) introduce EmotionCLIP, a self-supervised framework that captures emotional content from uncurated visual and verbal data by employing sentiment-guided contrastive learning and subject-aware encoding, illustrating the growing importance of multimodal emotional understanding.

Building on this foundation, an emerging literature in finance and accounting examines the informational role of nonverbal cues[1]. For example, Mayew and Venkatachalam (2012) demonstrate that vocal stress detected in managers' voices during earnings calls predicts future firm performance more accurately than verbal content alone. Similarly, Davila and Guasch (2021) use OpenPose[2] to measure entrepreneurs' body expansiveness during pitch presentations,

[1]Refer to Hanlon et al., 2022 for a comprehensive review of the literature.

[2]OpenPose is an algorithm developed by the CMU Perceptual Computing Lab. It uses a Unity plugin that jointly detects human body, hand, facial and foot keypoints on single images.

finding that expansive postures correlate with higher forecast errors, greater likelihood of funding success, but lower firm survival, indicating that nonverbal displays such as dominance and attractiveness materially influence investor perceptions.

Research in economics also leverages facial expression datasets to understand market dynamics. Breaban and Noussair (2018), using experimental data, find that traders' facial expressions of fear predict subsequent negative price movements, while positive expressions are linked to asset overpricing. In an empirical context, Alexopoulos et al. (2023) document that the Fed Chair's facial expressions during congressional testimony significantly impact financial markets, with effects comparable in magnitude to policy rate changes. These findings collectively underscore the powerful influence of emotional nonverbal cues on asset prices.

Curti and Kazinnik (2023) extend this line of inquiry by analyzing the Fed Chair's facial expressions during press conferences. They show that negative facial expressions correlate with negative stock returns even after controlling for the sentiment conveyed in words, suggesting that investors extract additional information from nonverbal signals. However, existing studies primarily focus on the correlation between emotional expressions and market outcomes, without investigating whether identical expressions are interpreted differently across Fed Chairs. This is a critical omission given that perceptions of credibility, experience, or familiarity may alter investors' emotional inferences.

Building on this gap, I develop the following hypothesis.

H1: Identical facial expressions elicit different market reactions depending on the Fed Chair

2.2. Investor Processing of nonverbal cues

Traditional Bayesian models assume that investors possess infinite memory, continuously and fully updating their beliefs based on incoming signals, such as the Fed Chair's expressions, regardless of the time elapsed since prior observations. Under this framework, each nonverbal cue would be weighted and processed as if investors retained perfect recall of the Chair's emotional history.

However, insights from psychology and behavioral finance suggest that information processing is more nuanced. According to the dual-processing framework, human cognition operates through two distinct systems: System 1, which is fast, automatic, and emotionally reactive, and System 2, which is slower, deliberate, and analytical (Kahneman, 2011). During salient public events such as FOMC press conferences, nonverbal cues may primarily activate System 1, prompting heuristic-based, emotionally charged belief revisions and resulting in immediate, pronounced market reactions.

Over time, as investors accumulate experience with a specific Fed Chair, their cognitive processing may shift. The bounded memory model suggests that the salience of emotional signals diminishes without reinforcement from new events (Wilson, 2014). Thus, the initial impact of a facial expression is likely to fade, reflecting a transition toward more System 2-driven analytical processing. Investors adapt by contextualizing the Chair's nonverbal cues within a broader informational framework, leading to more tempered market responses over time.

Over time, this results in a desensitization effect: as investors grow accustomed to a Fed Chair's typical expressions, negative emotional cues become less surprising and thus less impactful. Investors update their priors regarding the Chair's baseline expressiveness, leading to more muted market reactions. However, salient events such as congressional testimonies can act as attention resets—reactivating emotional salience and triggering heightened investor sensitivity. In such cases, investors may revert to fast, emotionally driven, System 1 processing. This incrases the market's sensitivity to nonverbal cues.

Integrating these insights, I propose the following hypotheses:

H2a: Investors react less negatively to negative facial expressions over time.

H2b: Investors' sensitivity to a Fed Chair's facial expressions increases following recent salient public events, such as congressional testimonies.

2.3. Facial Expressions' alignment with Fed's Goals

In psychology, Lambie and Marcel (2002) propose that individuals with higher emotional awareness develop more sophisticated emotion regulation strategies, such as cognitive reappraisal and suppression. Over time and with increased experience, individuals are theorized to become more adept at managing their emotional displays, especially in professional or high-stakes environments, where projecting stability and composure may be strategically important.

In the context of monetary policy communication, the Fed Chair serves as a central figure tasked with guiding market expectations and promoting financial stability. The introduction of post-meeting FOMC press conferences was intended to enhance transparency and predictability, thereby reducing policy uncertainty (Bernanke, 2013; Blinder, 1998; Amador and Weill, 2012). From a theoretical perspective, as Fed Chairs accumulate experience and become more familiar with the consequences of their public communications, one might expect them to improve their regulation of nonverbal signals—particularly by minimizing negative facial expressions that could undermine perceptions of confidence or stability.

Behavioral research underscores the influence of facial expressions on perceived trustworthiness and credibility. Positive emotional expressions, such as happiness, have been associated with enhanced perceptions of honesty and transparency (Stouten and De Cremer, 2010; Dotsch and Todorov, 2012; Hsieh et al., 2019). Conversely, negative emotional displays such as anger or contempt can reduce perceived trust and increase uncertainty. Given the Federal Reserve's longstanding objective of promoting market stability, there is a plausible expectation that Fed Chairs might, consciously or unconsciously, learn to reduce visible negative emotional signals over the course of their tenure. This is especially the case that deepfakes are realistic to the point that they are banned in Federal elections and also worldwide (Lynch, 2020; Chelvan and Chan, 2025).

While psychological research suggests that professional experience fosters improved emotional regulation, it is unclear whether this applies to the highly public and emotionally charged context of FOMC press conferences. Fed Chairs may learn to modulate their expressions, but the persistent pressures of financial market scrutiny and macroeconomic responsibility may override such regulation. It is possible that despite accumulated experience, the emotional pressures of high-stakes monetary policymaking continue to manifest in nonverbal behavior without significant suppression over time.

Accordingly, I propose the following hypothesis:

H3: Fed Chairs do not reduce their negative facial expressions during FOMC press conferences as their tenure increases.

This hypothesis tests whether emotional regulation, specifically the suppression of negative nonverbal cues, develops meaningfully over a Fed Chair's tenure, or whether facial expressions remain largely stable and unfiltered throughout their time in office.

H4: Investors react more strongly to facial expressions perceived as emotionally open (e.g., happy, neutral) than to emotionally opaque expressions (e.g., angry) during FOMC press conferences.

In this context, emotionally open expressions—such as happiness or calm—are defined as signals that investors associate with approachability, predictability, and transparency. These cues tend to lower perceived uncertainty. In contrast, emotionally opaque expressions like anger or contempt are less interpretable, more ambiguous in intent, and can amplify perceived risk, thereby triggering stronger market responses. For the rest of the paper, I adopt the "transparent" terminology to better reflect the conceptual link between emotional openness and perceived transparency in nonverbal communication.

3. Data and Methodology

3.1. Minute-level Market Data

To analyze how Federal Reserve Chairs' facial expressions during FOMC press conferences influence financial markets, I look at high-frequency market data capturing minute-by-minute movements in key asset prices. Using granular data allows for precise measurement of market reactions to nonverbal cues while minimizing the impact of other concurrent information releases, addressing potential endogeneity issues (Cochrane and Piazzesi, 2002; Nakamura and Steinsson, 2018).

I use the following variables to measure the market reaction to the nonverbal cues by the Chair during the FOMC press conference (from January 2011 to December 2020).

- SPDR S&P 500 (SPY): Minute-by-minute SPY prices and SPY trading volumes (number of shares traded).

- CBOE Volatility Index (VIX): Chicago Board Options Exchange Market Volatility Index (VIX), which measures implied volatility of the S&P 500.
- Euro-to-USD Exchange Rate (EUR): Minute-by-minute data for the Euro-to-USD exchange rate and its tick count per minute
- Japanese Yen-to-USD Exchange Rate (JPY): Minute-by-minute data for the JPY-to-USD exchange rate and its tick count per minute

For each asset, I calculate percentage changes within three-minute intervals during FOMC press conferences, measured in basis points. This interval aligns with the aggregation of facial expression data and facilitates a detailed examination of the immediate market response to nonverbal cues. Table 1 defines these variables, while Table 2 provides descriptive statistics.

3.2. Interpreting Nonverbal Cues with Advanced Facial Recognition Technology

To quantify the Fed Chairs' facial expressions during FOMC press conferences, I use advanced facial recognition technology. Specifically, I employ DeepFace, an open-source, state-of-the-art facial recognition and attribute analysis framework. DeepFace operates through a five-stage pipeline—detect, align, normalize, represent, and verify—to analyze facial expressions from video frames with high accuracy and scalability.

DeepFace wraps many state-of-the-art face recognition models - *VGG-Face , FaceNet, Open-Face, DeepFace, DeepID, ArcFace, Dlib, SFace and GhostFaceNet*. I use DeepFace v0.0.91, which is last updated in 2024 and the default configuation *VGG-Face model*. In technical terms, it uses a convolutional neural network (CNN) which represents faces as multi-dimensional vectors. DeepFace outputs emotional scores for seven facial emotions (*angry, fear, neutral, sad, disgust, happy and surprise*), where each emotion is scored from 0 to 100 and they sum up to 100. Fig. 2 provides an example of the measured frame.

To prepare the data, I collect FOMC press conference videos from the Federal Reserve's official channels, covering the period from April 2011 to December 2020. The videos are converted into frames at two-second intervals using a Python script. DeepFace is then applied to each frame to extract the emotional scores.

Recognizing that each Fed Chair's facial structure may influence how the software interprets their expressions, I account for this by calculating a baseline emotional profile for each Chair. This profile is obtained by averaging the emotional scores across all their press conferences, allowing for fair comparisons and controlling for individual facial characteristics.

To capture the dynamics of facial expressions over time, I aggregate the emotional scores into three-minute intervals, aligning with the market data intervals. The Negative Facial Expression Index is calculated as:

$$\text{Negative Facial} = \frac{\text{Angry}_{\text{3-mins average}} + \text{Fear}_{\text{3-mins average}} + \text{Disgust}_{\text{3-mins average}}}{\text{Angry}_{\text{chair lifetime average}} + \text{Fear}_{\text{chair lifetime average}} + \text{Disgust}_{\text{chair lifetime average}}}$$

Similarly, the Transparent Facial Expression Index is defined as:

$$\text{Transparent Facial} = \frac{\text{Happy}_{\text{3-mins average}} + \text{Neutral}_{\text{3-mins average}}}{\text{Happy}_{\text{chair lifetime average}} + \text{Neutral}_{\text{chair lifetime average}}}$$

These indices measure the relative intensity of negative and transparent expressions compared to each Chair's baseline, enabling assessment of deviations in emotional expressions during press conferences.

My initial sample includes 2657 minute-level observations from 46 FOMC press conferences held between April 2011 and September 2020. Of these, there are 18 press conferences with introductory statements. On average, each press conference lasts about 55 minutes, with the initial 10 minutes typically dedicated to the introductory statement.

Not all the screenshots are of the Fed Chair. To overcome this, I use a pre-trained VGG16 model from the Keras library. More technically, VGG16 is a 16-layer CNN model and the weights for the VGG16 model provided by the Keras library. Keras is a high-level neural networks API written in Python and integrated with TensorFlow. Keras is widely used in applied machine learning research due to its modular architecture, ease of use, and capacity to leverage pre-trained models such as VGG16, which has been trained on millions of images. This facilitates efficient image classification and feature extraction. I first manually locate a reference frame on the FOMC press conference of the Chair talking, check that this frame is similar in background with the majority of the frames of the Fed Chair, then use a cosine similarity test with a threshold of 50% to accurately sort the screenshots into whether they are of the Fed Chair or not.For example, there are frames which include a reporter asking questions or of a diagram being shown to explain FOMC policies. These are removed.

To ensure robustness, I only take screenshots where both the VGG16 cosine similarity test and the facial analysis displayed a N.A result. In other words, this checks for whether the Fed Chair is in the image or whether a face is on screen respectively, eliminating all non-Fed Chair frames. After processing, the final dataset comprises 1,440 minute-level observations from 46 FOMC press conferences. This dataset forms the basis for analyzing the relationship between the Fed Chairs' facial expressions and market reactions.

To further validate the reliability of the facial expression results obtained from DeepFace, I also use an independent facial emotion recognition system called **fer**, a popular open-source Python package first released in 2019. Unlike DeepFace, which is a general-purpose facial analysis framework, **fer** is purpose-built for recognizing emotions and has been trained on well-established emotion datasets. It detects faces in images and also classifies expressions into seven basic emotions: angry, disgust, fear, happy, sad, surprise, and neutral.

3.3. Controlling for Verbal Content Using Natural Language Processing (NLP)

To isolate the impact of nonverbal cues on financial markets, it is essential to control for the verbal content delivered during the FOMC press conferences. This ensures that observed market reactions can be accurately attributed to the Fed Chair's facial expressions rather than the verbal information conveyed.

To accurately synchronize between verbal content and video segments, I use **OpenAI Whisper**, a sophisticated speech recognition model developed by OpenAI. Whisper utilizes transformer architecture with an encoder-decoder structure, trained on extensive datasets that capture diverse accents and dialects. It effectively maintains context coherence over lengthy audio segments and provides highly accurate timestamps, ensuring that textual transcripts align closely with their corresponding visual frames. Additionally, to maintain accuracy and relevance, I remove segments where journalists ask questions during the Q&A, based on timestamps obtained from the earlier tests using the cosine similarity and VGG16-based classification, ensuring analysis is limited exclusively to the Fed Chair's speech.

To quantify and control for the informational content of verbal communication, I utilize a modified BERT-based model known as **FinBERT**, specifically fine-tuned for parsing financial text (Huang, Hui, and Yi, 2023). FinBERT systematically classifies speech segments into sentiment categories—positive, negative, and neutral—and further discerns nuanced monetary policy stances, identifying hawkishness or dovishness within the speech content. Additionally, I leverage **spaCy's** large language model (LLM) capabilities to extract linguistic features such as tokenization, keyword frequencies, entity recognition, and the identification of forward-looking statements. spaCy's advanced natural language processing tools complement FinBERT by providing a robust structural analysis of the text, allowing for a more granular and comprehensive control of verbal content.

For each analyzed three-minute interval, the transcribed text is processed using FinBERT in conjunction with spaCy LLM. This combined analysis yields comprehensive NLP measures, including sentiment scores, frequency counts of key financial and monetary policy keywords, hawkish/dovish indicators, and forward-looking statements. The resulting NLP features are explicitly selected to capture the multidimensional informational content relevant to monetary policy communication.

To ensure consistency and comparability across different press conferences, I standardize these NLP measures by normalizing them relative to their daily average:

$$\mathrm{NLP}_{k,t} = \frac{\mathrm{NLP}_{k,t'}}{\mathrm{NLP}_{k,\ \mathrm{day}}}$$

Here, k represents each NLP measure, for example, hawkishness, positivity, etc. t represents each minute, and t' represents each minute before averaging. For instance, to measure if the dialogue-sentence is positive or not, I use FinBERT to classify the sentence's sentiments. Since the financial figures and facial expressions are aggregated to each minute at the smallest level, I calculate a rolling-one-minute window of the count of the NLP and the average of all the rolling-one-minute counts in a day. I then take the count of the NLP measure divided by this average to derive the average score for each minute of FOMC press conference video. This procedure controls for variations in overall verbal content intensity across press conferences, thus enabling a precise isolation of facial expression effects.

3.4. Using Deepfake Simulations to Examine Investor Interpretations

A key contribution of this study is the use of deepfake simulations to investigate whether identical facial expressions elicit different market reactions depending on the Fed Chair. By creating deepfake videos that swap the faces of different Fed Chairs while preserving the underlying facial expressions (Li and Deng, 2022), I can isolate the effect of the Chair's identity on investor interpretations.

To create the deepfake videos, I use **DeepFaceLab**, an open-source framework for generating deepfakes. **DeepFaceLab** employs a combination of multi-task cascaded convolutional neural networks (CNNs), autoencoders, and other advanced algorithms for face swapping. The process involves several steps:

- **Face Detection and Alignment:** The software detects faces in the video frames and aligns them for consistent positioning.
- **Extraction and Preprocessing:** Facial features are extracted, and images are preprocessed to enhance quality.
- **Model Training:** An autoencoder model is trained to learn the facial features of both the source and target individuals.
- **Face Swapping and Post-processing:** The trained model swaps the faces in the video frames, ensuring that facial expressions and movements are preserved.

To select representative press conferences for each Fed Chair, I perform Principal Component Analysis (PCA) on the emotional intensity scores across all press conferences. The press conference closest to the centroid of each Chair's PCA plot is chosen, representing a typical emotional profile for that Chair. The selected press conferences are:

- Chair Ben Bernanke: March 2013 FOMC Press Conference
- Chair Janet Yellen: December 16, 2015 FOMC Press Conference
- Chair Jerome Powell: January 30, 2019 FOMC Press Conference

Using **DeepFaceLab**, I create deepfake videos where each Chair's face is swapped with another's while maintaining the original facial expressions and speech. For example, Chair Powell's face is superimposed onto Chair Yellen's body during her press conference. The deepfake models are trained extensively, up to 100,000 iterations, to ensure high-quality and realistic videos. More details can be found in Appendix A4.

Figure 1 shows a deepfake model, which overlays Chair Powell's face on Chair Yellen's. These deepfake simulations provide a novel methodological approach to testing Hypothesis H1, which posits that identical facial expressions elicit different market reactions depending on the Fed Chair.

To ensure that these results are robust and replicable, I independently replicate the face-swapping pipeline using **faceswap.py**, another open-source deepfake implementation that integrates with TensorFlow. Unlike DeepFaceLab, which is optimized for high customizability and fine-grained training workflows, **faceswap.py** provides a user-friendly interface and built-in GPU acceleration via TensorFlow backends. The secondary implementation allows me to verify that the facial integrity and emotional consistency of the swapped videos are preserved regardless of the specific software used—adding credibility to the main identification strategy.

3.5. Other Control Variables

To account for other factors that might influence market reactions during FOMC press conferences, I incorporate several control variables in the analysis. These variables help isolate the specific effect of the Fed Chair's facial expressions on financial markets:

- **Federal Funds Rate Change (ΔFFR):** The change in the target federal funds rate announced during the FOMC meeting. Controlling for ΔFFR accounts for the direct impact of monetary policy decisions on market movements.
- **Pre-FOMC Drift Variables:** To control for the documented pre-FOMC announcement drift in asset prices (Lucca and Moench, 2015; Boguth et al., 2019), I include variables representing the percentage change in SPY, VIX, EUR/USD, and JPY/USD prices within the 30 minutes before the press conference.
- **Monetary Policy Uncertainty (MPU) Index:** The MPU index measures public uncertainty about Federal Reserve policy actions and their consequences (Husted et al., 2020). Including the MPU index controls for the general level of uncertainty surrounding each FOMC meeting.

- **Online Attention Measures:** To capture the level of public and investor attention to each FOMC press conference, I include variables such as Google Trends data related to the Fed or monetary policy on the day of the press conference.

4. Results

4.1 Are the same expressions by different Fed Chairs read differently?

In my preliminary result, I first ask whether facial expressions are a complex signal. Table 3 shows the min-max normalized changes of analyzing a twenty-second video of each Fed Chair that has their faces changed via the deepfake technique. Each deepfake takes approximately four hours to train with a given high-spec GPU stated in Appendix A4. The results indicate that even though the underlying expression is the same, when the face is changed, the facial emotion recognition software registers different emotions. *disgust* changes the most between Yellen's original face and overlaying it with Powell's face. Across the board, we see that *sad*, *surprise* and *neutral* are the most stable while the other emotions change more widely.

This suggests that there is a baseline of emotions for each Fed Chair and that facial expressions are not straightforward to interpret across individuals. I posit that this lends support to the macroeconomic difference-in-opinion model proposed by Allen, Morris, and Shin (2006) and Banerjee et al. (2009), which emphasizes how public signals can be subject to varying interpretations among investors. In this setting, facial expressions function as such public signals —but the finding that identical expressions elicit different emotional readings depending on the face presenting them implies that investors may interpret these cues through the lens of prior beliefs about the Chair's personality, credibility, or communication style. Rather than processing facial expressions in isolation, investors appear to weigh them against contextual expectations. This variation in interpretation aligns with the idea that heterogeneous beliefs can emerge even when all investors receive the same observable signal.

Untabulated results continue to support these results. Using an independent set of deepfakes generated via faceswap.py, I then used fer, which is a different emotion recognition package.

Even when facial movements are held constant, the perceived emotional content differs systematically depending on the identity of the Fed Chair.

This observation leads me naturally to further investigate what kind of cognitive processing model investors use when interpreting nonverbal cues in Section 4.6.

4.2 Do investors react to negative facial expressions?

I next examine whether Chairs' emotions are related to the changes in stock and currency market. I use a fixed-effects regression model and estimate the following model:

$$\%\Delta\text{Market}_{\text{t}} = \alpha + \beta_1 \text{Negative Facial}_{\text{t-1}} + \text{controls} + \varepsilon_{\text{t}}$$

For each market reaction $\%\Delta\text{Market}_{\text{t}}$, I calculate the absolute percentage change of each 1-minute interval In untabulated results, I find no meaningful results if I use percentage change without the absolute sign. My results replicate Curti and Kazinnik (2023) as they find that expressions cause reactions from the market. I take this to provide robustness to my data collection and methods. If the Fed Chair makes a negative expression, investors may perceive it as good and bad depending on how they interpret it. For example, when Chair Powell says *"We continue to anticipate that ongoing increases in the target range for the federal funds rate will be appropriate in order to attain a stance of monetary policy that is sufficiently restrictive."*, a negative expression could be interpreted as federal fund rates increasing due to the vibrant business economy, which would be interpreted as bullish for the markets. It could also be interpreted as the Fed's increase have negative impact on borrowing and thus, causing the market to become bearish. The results in Table 4 show that negative facial expressions are significantly associated with higher market volatility, particularly for the VIX and currency markets (columns 3–5 and 8–9), suggesting that investors react to visual emotional cues even after controlling for verbal sentiment, policy stance, and other relevant variables. The negative coefficients on the "Negative Facial" variable imply that greater perceived negativity from the Chair increases uncertainty or perceived risk, leading to larger absolute price changes.

To further investigate the relationship between the markets and the Fed Chair's facial expressions, I look at trading volume and tick count. I perform a multivariate regression similar to the previous regression in this sub-section, but with the dependent variable changed to volume.

Table 5 corroborates this interpretation using volume-based measures. Negative facial expressions significantly predict higher stock volume and VIX tick activity, reinforcing the idea that investors not only respond emotionally but also actively trade in response to these nonverbal cues. The consistent significance across multiple asset classes and measures of market activity suggest that facial expressions are not simply noise—they are perceived as meaningful signals in financial markets.

4.3 Does the Fed Chair align her facial expressions with Fed's goal of stability?

If facial expressions indeed affect investors' assessment, then the next relevant question would be whether the Fed Chair consciously controls their facial expressions. If the Fed Chair knew that negative reactions cause instantaneous reactions from the market, they should decrease their use of negative expressions while increasing the frequency of other expressions such as neutral or happy.

Table 6, shows that as the Fed Chair tenure increases, the frequency and intensity of their negative expressions increase. Columns (1) and (2) show that conference count directly increases the negative expressions. From this result, I conjecture that the Fed Chairs do not strategically control facial expressions. Since more negative expressions elicit more negative reactions (Curti and Kazinnik, 2023), and that it also goes against the goal of reducing market volatility, Fed Chairs should reduce such expressions if they were strategic[3]. In fact, this problem was highlighted in that excessive Federal Reserve communication leads to much undue volatility (El-Erian, 2023). I bear in mind that FOMC press conferences previously under Bernanke and Yellen were held only on select dates but after 2018, they are held eight times a year. Whether this volatility is a conscious effort of the Fed Chair cannot be explicitly tested but in line with the implicit goal of non-volatility, I conclude that there is no strategic use on the Fed Chair's part.

In columns (4) - (7), I show that as the conference count increases, the Fed Chair also decreases the ratio of neutral expressions. The Fed Chair shows more negative expressions and

[3]The results are robust to varying economic conditions. It was pointed out that this result could be because of the age of the Fed Chair causing more negative expressions (Grimmer et al., 2021). However, in this dataset, the maximum tenure is four years for Yellen. It is unlikely that they massively aged in this timespan.

less neutral expressions, while there is no statistically significant relationship for happy or sad expressions.

These findings in Table 6 provide suggestive evidence against the notion that Fed Chairs consciously calibrate their facial expressions to influence market outcomes. If facial expressions were used as strategic communication tools, one might expect a decline in negative emotional displays with experience—particularly given the well-documented sensitivity of markets to negative cues. However, the data shows the opposite. As Fed Chairs gain experience (measured by conference count), the frequency of negative expressions increases, while the proportion of neutral expressions declines.

This pattern implies that emotional displays during press conferences may reflect authentic cognitive or emotional states rather than being deliberately controlled for signaling purposes. The absence of significant changes in happy or sad expressions further supports the idea that these cues are not finely tuned for market interpretation. Together, these findings point toward a more human and unscripted dimension of monetary policy communication.

4.4 Do investors react to facial expressions that convey transparency?

Building on the literature that focuses primarily on negative expressions or sentiments to quantify FOMC press conferences, I extend the analysis to investigate whether facial expressions associated with transparency, such as happiness and neutrality, also influence market outcomes. These expressions, which I earlier termed as transparent, are theorized to signal emotional openness and communicative clarity, which could potentially affect investors' perceptions

Table 7, Columns (1) through (9), examine the effects of non-negative expressions on various market indicators. Statistically significant coefficients in Columns (3), (6), and (7) suggest that transparent expressions do elicit market responses, albeit not as consistently as negative expressions. This supports H4, which posits that emotionally open cues—such as happiness and calm—can influence investor behavior during press conferences.

From an information economics perspective, transparent emotional expressions may be perceived by investors as signals of a Chair's willingness to communicate honestly and credibly. Such expressions could reduce perceived uncertainty and enhance trust in the Fed's policy intentions. Behavioral research lends support to this interpretation: prior studies link positive affect displays to heightened perceptions of credibility, stability, and honesty (Stouten and De Cremer, 2010; Dotsch and Todorov, 2012).

Interestingly, while individual expressions such as happiness or neutrality are only occasionally significant, the broader measure of transparent emotional expression emerges as a more robust predictor of market responses. This suggests that investors may interpret emotional cues not in isolation, but as part of a higher-order emotional composite or communicative context. In doing so, they move beyond the seven basic emotions and toward a more nuanced interpretation of emotional tone, especially under high-stakes settings like monetary policy announcements.

These findings reinforce the idea that nonverbal signals, including subtle affective displays, form an important component of monetary policy communication. They further support the theoretical view that investor reactions are not solely tied to the content of verbal guidance, but are shaped by perceived transparency and emotional openness as conveyed through facial expressions.

4.5 Do investors react to facial expressions that contrast with word sentiment?

While earlier sections show that investors react to negative facial expressions, an important question arises - Is this effect driven by emotional dissonance, defined as when a Fed Chair's facial expression contradicts the sentiment of their spoken words? If investors are particularly attuned to such inconsistencies, then facial expressions that diverge from verbal sentiment (e.g., a frown during optimistic guidance) might intensify market reactions.

To test this, I interact negative sentiment scores (measured using FinBERT) with negative facial expression indices in a fixed-effects regression framework. Table 8 presents the results. The coefficients on the interaction terms are statistically insignificant across all specifications (Columns 1–5), suggesting that markets do not respond more strongly when facial expressions and verbal sentiment diverge. In other words, there is no evidence that emotional dissonance amplifies investor reactions.

Instead, I interpret these findings as suggesting that facial expressions function as an independent communicative signal. Investors appear to process nonverbal cues on their own, rather than only in conjunction with the verbal content. This supports a dual-channel interpretation of central bank communication, where expressions and words are evaluated separately. It also aligns with long-standing insights from nonverbal communication research (e.g., Darwin, 1872; Mehrabian, 1972), which suggest that facial expressions can convey meaning even when not aligned with speech.

Taken together, the results imply that investors do not necessarily punish or reward inconsistency between words and expressions, but instead, interpret facial expressions as standalone cues about policy intentions, emotional conviction, or uncertainty.

4.6 How do investors differentially interpret negative facial expressions?

Investor interpretation of facial expressions may evolve over time as the Fed Chair gains experience. Curti and Kazinnik (2023) document that negative expressions are increasingly well understood by markets as the Chair's tenure progresses. I extend this analysis by examining whether market responses to negative facial expressions are moderated by the Chair's cumulative experience and recent events.

Table 9 reports results from interaction regressions that include both Fed Chair tenure and key event-based controls. Columns (1), (2), (4), and (6) show that the interaction between tenure and negative facial expressions is statistically significant. This finding corroborates Curti and Kazinnik (2023), indicating that investors gradually learn how to interpret each Chair's nonverbal cues. Interestingly, the coefficients in Columns (2) and (6) are negative, suggesting that the absolute magnitude of market reactions declines with tenure. This is consistent with an experience-based learning mechanism: over time, markets may become more confident in interpreting a given Chair's expressions, resulting in lower volatility and more muted responses.

In Columns (5) and (8), I condition on whether a congressional testimony occurred shortly before the FOMC press conference. The interaction term is positive and significant, indicating that in these instances, negative facial expressions have a stronger effect on market volatility. I interpret this as evidence of context-sensitive interpretation. Congressional testimonies may impose heightened emotional or reputational pressure, leading to expressions that investors perceive as more salient or revealing. These events can Chair's usual communication pattern and temporarily amplify investor sensitivity to nonverbal cues.

Furthermore, Columns (3) and (7) introduce a career progression variable using quartiles of the Chair's term. The results indicate that as Chairs approach later stages of their tenure, negative expressions exert less influence on market responses. This suggests an adaptive investor belief formation process—markets update expectations over time, potentially discounting expressions viewed as routine or less informative in later tenures.

Taken together, these results support a bounded-memory dual-processing model. Investors appear to form learned interpretations of each Chair's facial cues, with reactions dampening over time as familiarity increases. However, memory is not infinite because recent events such as testimonies serve as exogenous shocks that recalibrate investor attention and restore sensitivity to facial signals. This dynamic adjustment process aligns with models of market learning under informational frictions (e.g., Wilson, 2014), where investor interpretations evolve but remain susceptible to context-specific disruptions.

5 Conclusion

This paper introduces a new channel through which central bank communication affects financial markets: the facial expressions of the Fed Chair. While existing studies focus predominantly on the verbal content of FOMC press conferences, I show that nonverbal cues—specifically facial expressions—represent a distinct and economically meaningful signal that shapes investor reactions in real time.

A key methodological contribution of this study is the use of deepfake simulations to isolate the interpretive role of the Fed Chair's identity. By holding facial expressions constant and altering only the face using high-resolution deepfake videos, I demonstrate that identical emotional expressions are interpreted differently depending on who is delivering them. This finding implies that facial expressions are not read in isolation but are filtered through the lens of investor priors—such as perceived credibility, communicative style, or past behavior. These results lend support to the theoretical foundation of difference-in-opinion models, where public signals take on divergent meanings depending on who conveys them.

Investor responses to facial expressions also appear to evolve over time. As investors become more familiar with a particular Fed Chair, the market's sensitivity to their negative expressions diminishes. This suggests a learning mechanism through which investors recalibrate their interpretations, reducing the volatility associated with familiar emotional signals. However, this learning is not unbounded: salient recent events—such as congressional testimonies—reactivate investor attention and lead to stronger responses. These findings are consistent with a dual-processing and bounded-memory framework, in which emotionally salient signals decay in influence but can be reweighted by new contextual cues.

Importantly, the paper also finds that facial expressions interpreted as transparent, such as happiness or neutrality, can be equally, if not more, impactful in certain market contexts. These expressions may signal approachability or clarity in communication, prompting investors to revise expectations about policy certainty. This broadens the traditional focus on negative emotional cues by highlighting the market's sensitivity to emotional openness as a proxy for transparency.

Finally, I examine whether Fed Chairs strategically regulate their facial expressions to manage market reactions. Contrary to theoretical expectations rooted in emotional regulation literature, the evidence suggests that facial expressions become more—not less—negative over the course of a Chair's tenure. This points away from intentional control and suggests that the emotional toll of high-stakes monetary policymaking continues to manifest in nonverbal behavior.

Several limitations are acknowledged. FER systems are sensitive to video quality, model architecture, and underlying training data. While I use two distinct systems, DeepFace and **fer**, to triangulate emotional measurements, future experimental approaches may provide cleaner identification. Moreover, the 3-minute aggregation window used in the analysis may smooth over short-term variations in expressions, though this is necessary to align with financial market data frequency.

Another key limitation of this study is that the deepfake counterfactuals were not shown to market participants directly. As such, the analysis relies on observed market reactions to real press conferences, while the counterfactuals are used to estimate the differential effects of nonverbal cues ex post. This implies that while the design permits causal interpretation under strong identification assumptions, it does not measure investor reactions to the counterfactual content in a behavioral or experimental setting. Future research could complement this approach with survey or lab-based methods that directly expose participants to the manipulated videos.

Taken together, this study introduces a novel methodological and conceptual framework for understanding nonverbal communication in monetary policy. By combining deepfake simulations, facial emotion recognition, and natural language processing, the paper opens new avenues for research on how central banks influence expectations—not only through what is said, but through how it is expressed.

References

Alexopoulos, M., Han, X., Kryvtsov, O., & Zhang, X. (2024). More than words: Fed Chairs' communication during congressional testimonies. Journal of Monetary Economics, 142, 103515.

Allen, F., Morris, S., & Shin, H. S. (2006). Beauty contests and iterated expectations in asset markets. The Review of Financial Studies, 19(3), 719–752.

Amador, M., & Weill, P.-O. (2012). Learning from private and public observations of others' actions. Journal of Economic Theory, 147(3), 910–940.

Amengual, D., & Xiu, D. (2018). Resolution of policy uncertainty and sudden declines in volatility. Journal of Econometrics, 203(2), 297–315.

Banerjee, S., Kaniel, R., & Kremer, I. (2009). Price drift as an outcome of differences in higher-order beliefs. The Review of Financial Studies, 22(9), 3707–3734.

Bernanke, B. (2007). Federal Reserve Communications. Speech at the Cato Institute 25th Annual Monetary Conference, Washington, DC.

Bernanke, B. (2013). A Century of US Central Banking: Goals, Frameworks, Accountability. Journal of Economic Perspectives, 27(4), 3–16. https://doi.org/10.1257/jep.27.4.3

Boguth, O., Grégoire, V., & Martineau, C. (2019). Shaping Expectations and Coordinating Attention: The Unintended Consequences of FOMC Press Conferences. Journal of Financial and Quantitative Analysis, 54(6), 2327–2353. https://doi.org/10.1017/S0022109018001357

Breaban, A., & Noussair, C. N. (2018). Emotional state and market behavior. Review of Finance, 22(1), 279–309.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. Journal of Monetary Economics, 12(3), 383–398.

Chelvan, V., & Chan, G. (2025). Law banning deepfakes of candidates in place for first time in GE2025, among other guidelines: ELD | The Straits Times. The Straits Times. https://www.straitstimes.com/singapore/politics/new-law-banning-fake-online-ads-to-be-in-place-for-first-time-in-ge2025-among-other-guidelines-eld

Cochrane, J. H., & Piazzesi, M. (2002). The Fed and Interest Rates—A High-Frequency Identification. American Economic Review, 92(2), 90–95. https://doi.org/10.1257/000282802320189069

Coibion, O., Gorodnichenko, Y., & Kamdar, R. (2018). The formation of expectations, inflation, and the phillips curve. Journal of Economic Literature, 56(4), 1447–1491.

Curti, F., & Kazinnik, S. (2023). Let's face it: Quantifying the impact of nonverbal communication in FOMC press conferences. Journal of Monetary Economics, 139, 110–126. https://doi.org/10.1016/j.jmoneco.2023.06.007

Darwin, C. (1872). The Expression of Emotion in Man and Animals, by Charles Darwin. https://www.gutenberg.org/files/1227/1227-h/1227-h.htm

Davila, A., & Guasch, M. (2022). Managers' Body Expansiveness, Investor Perceptions, and Firm Forecast Errors and Valuation. Journal of Accounting Research, 60(2), 517–563. https://doi.org/10.1111/1475-679X.12426

Dotsch, R., & Todorov, A. (2012). Reverse Correlating Social Face Perception. Social Psychological and Personality Science, 3(5), 562–571. https://doi.org/10.1177/1948550611430272

El-Erian, M. (2023, November 20). Federal Reserve Officials Talk Too Much. Bloomberg.Com. https://www.bloomberg.com/opinion/articles/2023-11-20/federal-reserve-officials-talk-too-much

Emett, S. A., Eulerich, M., Pickerd, J. S., & Wood, D. A. (2024). Short and Synthetically Distort: Investor Reactions to Deepfake Financial News. Available at SSRN 4869830. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4869830

Geraats, P. M. (2009). Trends in Monetary Policy Transparency. International Finance, 12(2), 235–268. https://doi.org/10.1111/j.1468-2362.2009.01239.x

Gómez-Cram, R., & Grotteria, M. (2022). Real-time price discovery via verbal communication: Method and application to Fedspeak. Journal of Financial Economics, 143(3), 993–1025. https://doi.org/10.1016/j.jfineco.2021.12.004

Gorodnichenko, Y., Pham, T., & Talavera, O. (2023). The voice of monetary policy. American Economic Review, 113(2), 548–584.

Grimmer, M., Ramachandra, R., & Busch, C. (2021). Deep Face Age Progression: A Survey. IEEE Access, 9, 83376–83393. IEEE Access. https://doi.org/10.1109/ACCESS.2021.3085835

Hanlon, M., Yeung, K., & Zuo, L. (2022). Behavioral Economics of Accounting: A Review of Archival Research on Individual Decision Makers. Contemporary Accounting Research, 39(2), 1150–1214. https://doi.org/10.1111/1911-3846.12739

Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A Large Language Model for Extracting Information from Financial Text. Contemporary Accounting Research, 40(2), 806–841. https://doi.org/10.1111/1911-3846.12832

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux. https://www.pdcnet.org/pdc/bvdb.nsf/showopenaccess?open&repid=852577BA0050DC0D&docid=2BE9C7FA6C507E4DC1257ADA0072A186&solarid=inquiryct_2012_0027_0002_0054_0057

Kaur, J., Saxena, J., Shah, J., & Yadav, S. P. (2022). Facial Emotion Recognition. 2022 International Conference on Computational Intelligence and Sustainable Engineering Solutions (CISES), 528–533. https://ieeexplore.ieee.org/abstract/document/9844366/?casa_token=Ibhp_QgNdUMAAAAA:1itJ3wCR1_CZW3Y9OVeMMuTl_OYPk6bS9iXxGjGsLdkrqN-x1yjeOjTCtEHjVMV-xT4xGSDNecjm

Lambie, J. A., & Marcel, A. J. (2002). Consciousness and the varieties of emotion experience: A theoretical framework. Psychological Review, 109(2), 219–259. https://doi.org/10.1037/0033-295X.109.2.219

Li, S., & Deng, W. (2022). Deep Facial Expression Recognition: A Survey. IEEE Transactions on Affective Computing, 13(3), 1195–1215. https://doi.org/10.1109/TAFFC.2020.2981446

Lucas, R. E. (1972). Expectations and the Neutrality of Money. Journal of Economic Theory, 4(2), 103–124.

Lucca, D. O., & Moench, E. (2015). The Pre-FOMC Announcement Drift. The Journal of Finance, 70(1), 329–371. https://doi.org/10.1111/jofi.12196

Mayew, W. J., & Venkatachalam, M. (2012). The Power of Voice: Managerial Affective States and Future Firm Performance. The Journal of Finance, 67(1), 1–43. https://doi.org/10.1111/j.1540-6261.2011.01705.x

Mehrabian, A. (1972). Nonverbal Communication. Transaction Publishers.

Nakamura, E., & Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. The Quarterly Journal of Economics, 133(3), 1283–1330.

Neuhierl, A., & Weber, M. (2019). Monetary policy communication, policy slope, and the stock market. Journal of Monetary Economics, 108, 140–155. https://doi.org/10.1016/j.jmoneco.2019.08.005

Renier, L. A., Shubham, K., Vijay, R. S., Mishra, S. S., Kleinlogel, E. P., Jayagopi, D. B., & Schmid Mast, M. (2024). A deepfake-based study on facial expressiveness and social outcomes. Scientific Reports, 14(1), 3642. https://doi.org/10.1038/s41598-024-53475-5

Lynch, S. F. [D-M.-8. (2020, March 4). H.R.6088 - 116th Congress (2019-2020): Deepfakes in Federal Elections Prohibition Act (2020-03-04) [Legislation]. https://www.congress.gov/bill/116th-congress/house-bill/6088

Stouten, J., & De Cremer, D. (2010). “Seeing is believing”: The effects of facial expressions of emotion and verbal communication in social dilemmas. Journal of Behavioral Decision Making, 23(3), 271–287. https://doi.org/10.1002/bdm.659

Swanson, E., & Jayawickrema, V. (2023). Speeches by the fed chair are more important than fomc announcements: An improved high-frequency measure of us monetary policy shocks. Unpublished Manuscript. https://sites.socsci.uci.edu/~swanson2/papers/hfdat.pdf

Westerlund, M. (2019). The Emergence of Deepfake Technology: A Review. 39–52.

Wilson, A. (2014). Bounded Memory and Biases in Information Processing: Bounded Memory and Biases in Information Processing. Econometrica, 82(6), 2257–2294. https://doi.org/10.3982/ECTA12188

Woodford, M. (2001). Monetary Policy in the Information Economy (Working Paper 8674). National Bureau of Economic Research. https://doi.org/10.3386/w8674

Yellen, J. L. (2016). Opening remarks: The Federal Reserve's monetary policy toolkit: Past, present and future. Designing Resilient Monetary Policy Frameworks for the Future, 1–19.

Tables and Figures

Panel A: Fed Chair Yellen speaking during a FOMC press conference. Dominant emotion identified by facial analysis software - **Sad**. Complete Analysis: Angry: 0.722%, Disgust: 0.036%, Fear: 21.992%, Happy: 0.057%, Sad: 58.435%, Surprise: 0.021%, Neutral: 18.737%

Panel B: Deepfake of Fed Chair Yellen speaking, using Fed Chair Powell's face, during a FOMC press conference. Dominant emotion identified by facial analysis software - **Angry**. Complete Analysis: Angry: 99.957%, Disgust: 0.0000038%, Fear: 0.001%, Happy: 0.000%, Sad: 0.041%, Surprise: 0.000%, Neutral: 0.000%

Figure 1. Comparison of a deepfake of Fed Chair Janet Yellen during FOMC press conference on September 21, 2016 using Fed Chair Powell and their facial analysis result

Panel A: Ben Bernanke, April 27 2011	
Emotion	Score
Angry	21.0
Disgust	0.0
Fear	0.2
Happy	60.0
Sad	12.7
Surprise	0.0
Neutral	5.8

Panel B: Janet Yellen, June 14, 2017	
Emotion	Score
Angry	47.0
Disgust	0
Fear	2.6
Happy	10.2
Sad	19.3
Surprise	0.2
Neutral	20.7

Panel C: Jerome Powell, Sept 26 2018	
Emotion	Score
Angry	0.20
Disgust	0
Fear	0
Happy	0.1
Sad	2.5
Surprise	0
Neutral	97.2

Figure 2. Emotion Scores: The emotion intensity scores are captured by DeepFace. Panel A is Ben Bernanke during the FOMC press conference on April 27th, 2011. Panel B is Janet Yellen during the FOMC press conference on June 14th, 2017. Panel C is Jerome Powell during the FOMC press conference on September 26th, 2019.

Table 1

Variable Definitions. This table presents definitions of the variables used in the paper.

Variables	Definition	Source
%Δ SPY	Percent change in SPY (SPDR S&P 500), measured every minute	Bloomberg
%Δ VIX	Percent change in VIX (CBOE Volatility Index), measured every minute	Bloomberg
%Δ EUR	Percent change in spot EUR-USD exchange rate, measured every minute	Bloomberg
%Δ JPY	Percent change in spot JPY-USD exchange rate, measured every minute	Bloomberg
SPY Volume	SPY trading volume in a 1-minute interval	Bloomberg
EURUSD Tick Count	EURUSD number of tick counts in a 1-minute interval	Bloomberg
JPYUSD Tick Count	EURUSD number of tick counts in a 1-minute interval	Bloomberg
Independent Variables		
Negative Facial	Chair's intensity of transparent facial expressions averaged in the prior three minutes divided by average transparent facial expressions across all FOMC meetings by the Chair	DeepFace
Transparent Facial	Chair's intensity of transparent facial expressions averaged in the prior three minutes divided by average transparent facial expressions across all FOMC meetings by the Chair	DeepFace
Neutral Facial	Chair's intensity of transparent facial expressions averaged in the prior three minutes divided by average transparent facial expressions across all FOMC meetings by the Chair	DeepFace
Happy Facial	Chair's intensity of transparent facial expressions averaged in the prior three minutes divided by average transparent facial expressions across all FOMC meetings by the Chair	DeepFace
Sad Facial	Chair's intensity of transparent facial expressions averaged in the prior three minutes divided by average transparent facial expressions across all FOMC meetings by the Chair	DeepFace
predrift_k	Percent change in *k* from 2.00pm to 2.30pm, the time between when FOMC statement is released and FOMC press conference is held, where *k* is SPY, VIX, EURUSD, or JPYUSD.	Bloomberg

Table 1

Variable Definitions. This table presents definitions of the variables used in the paper.

Variables	Definition	Source
cfquart	Dummy variable of a chair's FOMC press conference throughout the four quartiles of their count of FOMC press conference	
congre30	Dummy variable of whether a congressional testimony was held within 30 days prior to the FOMC press conference	
congre10	Dummy variable of whether a congressional testimony was held within 10 days prior to the FOMC press conference	
%Δ FDFD	Percent change in ICAP US Federal Funds Rate Index on day of FOMC announcement, measured daily	Bloomberg
MPU	Monetary Policy Uncertainty (MPU), measured daily	Bloomberg
Public_Interest	3-day average of Google Search Index Value before the actual date of FOMC Press Conference. Recommended Keywords by Google most related to the topic - "FOMC Meeting"	Google
NLP Variables		
Negative Sentiment	FinBERT Measure of negativity of a statement. FinBERT is a Large Language Model specialized for financial language. (Huang et al. 2022)	
Statement-related	Measures whether the statement is related to the FOMC-statement released at 2.00pm (Gomez-Cram and Grotteria, 2021) I use spaCy, a English language model designed for NLP tasks to tokenize the corpus and search against keywords	
Hawkish	Measures in binary format whether statement is hawkish or dovish using word-match (Neuhierl and Weber, 2019) I use spaCy, a English language model designed for NLP tasks to tokenize the corpus and search against keywords	

Table 1

Variable Definitions. This table presents definitions of the variables used in the paper.

Variables	Definition	Source
FLS	Measures how forward-looking each statement is based on a list of key words. I use spaCy, a English language model designed for NLP tasks to tokenize the corpus and search against keywords	

Table 2

Descriptive Statistics

Variable	Count	Mean	Std Dev	Min	25%	50%	75%	Max
%Δ SPY	1471	0.039	0.043	0	0.012	0.028	0.050	0.538
%Δ VIX	1471	0.272	0.391	0	0.065	0.151	0.341	4.509
%Δ EUR	1471	0.025	0.029	0	0.008	0.017	0.035	0.356
%Δ JPY	1471	0.023	0.025	0	0.009	0.018	0.031	0.259
Volume SPY	1456	0.478	0.502	0.027	0.179	0.339	0.629	11.083
Tick_Count_vix	1471	4.750	1.564	3.000	4.000	4.000	4.000	8.000
EUR Tick	1471	1.919	1.427	0.247	0.982	1.480	2.379	10.378
Negative Facial	1465	1.206	0.481	0.185	0.858	1.151	1.470	3.083
Transparent Facial	1465	1.786	0.849	0.222	1.200	1.605	2.157	4.454
Neutral Facial	1465	1.208	0.790	0.001	0.644	1.050	1.600	5.654
Happy Facial	1465	1.233	1.203	0.000	0.417	0.900	1.679	9.860
Sad Facial	1465	1.194	0.766	0.004	0.648	1.058	1.577	5.385
Predrift SPY	1456	0.126	0.408	−1.020	−0.163	0.093	0.366	1.060
Predrift VIX	1426	−1.958	3.431	−13.009	−3.226	−1.220	0	4.618
Predrift EUR	1471	0.059	0.325	−0.874	−0.042	0.072	0.156	0.920
Predrift JPY	1471	0.020	0.317	−0.836	−0.131	−0.038	0.196	1.017
Negative Sentiment	1471	0.700	0.726	0	0	0.526	1.138	3.699
Statement_Related	1410	1.134	3.729	0	0	0	0	64.000
%Δ FDFDperChange	1471	0.030	0.142	−0.417	0	0	0.100	0.610

Hawkish	1471	0.552	5.028	0	0	0	0	66.500
Public_Interest	1437	51.179	9.243	33.333	45.000	51.667	58.333	66.667

Table 3. Are the same expressions by different Fed Chairs read differently? This table shows the min-max normalized changes when comparing coefficients using 20-seconds of deepfake video trained on an average of 100,000 iterations of each of the 7 basic emotions. The average emotions displayed through the 20 seconds are derived from 2-second interval screenshots. Using the deepfake technology, the face structure changes but the underlying emotions are the same. We see that fear and surprises displays the lowest changes while disgust shows a much larger variation. Data points highlighted in **bold** show the maximum and minimum differences.

Emotion	(1) Yellen Original vs Yellen Original Overlaid with Powell's Face	(2) Powell Original vs Powell Original Overlaid with Bernanke's Face	(3) Bernanke Original vs Bernanke Original Overlaid with Yellen's Face
angry	0.258	0.008	0.003
disgust	**1.000**	0.207	0.215
fear	**0.000**	0.190	0.008
happy	0.004	0.068	0.242
sad	0.018	0.005	0.017
surprise	0.002	0.014	**0.000**
neutral	0.007	0.013	0.011

Table 4. Do investors react to negative facial expressions? This table presents coefficients from OLS regressions examining changes in stock (SPY), currency (EUR), (JPY), and the VIX volatility index in response to FOMC chairs' negative emotions and control variables. The analysis includes 1359 to 1404 observations at the minute level spanning 46 FOMC meetings chaired by Ben Bernanke (12), Janet Yellen (16), and Jerome Powell (18) from April 27th, 2011, to September 16th, 2020. Percent changes in SPY, VIX, EUR, JPY are measured over each minute and the absolute value is taken. Negative Facial Expressions represents the intensity of chairs' negative emotions averaged over the preceding three minutes relative to the average across all meetings under the chair. This is to control for the specific nature and disposition of each Fed Chair. Negative Sentiment measures the expressed tone based on FinBERT for each statement . Hawkishness measures the policy stance of chairs based on the keyword list in (Neuhierl and Webet, 2019) and spaCy LLM tokenization. Statement Related measures the frequency of statements in a time interval that are related to the FOMC Press Statement given at 2.00pm. All language parameters are averaged over each rolling minute. Predrift captures percent changes in the 30 minutes from 2.00pm to 2.30pm before the FOMC press conference for SPY, VIX, EUR, JPY respectively. MPU indicates the Monetary Policy Uncertainty index before the FOMC meeting as per Husted et al. (2020). %ΔFDFD denotes the change in Federal Funds Rate on the day of FOMC Press Conference. Standard errors, shown in parentheses, are clustered at the FOMC meeting level. Variable definitions are detailed in Table 1.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	%Δ SPY	%Δ SPY	%Δ SPY	%Δ VIX	%Δ VIX	%Δ VIX	%Δ EUR	%Δ EUR	%Δ JPY
Negative Facial	−0.007	−0.008	−0.007**	−0.016	−0.083**	−0.065*	0.003	0.001	0.005**
	(0.005)	(0.006)	(0.002)	(0.043)	(0.040)	(0.033)	(0.003)	(0.003)	(0.002)
Negative Sentiment	0.003*	0.003*	0.003*	0.010	0.009	0.013	0.001	0.001	0.002
	(0.001)	(0.001)	(0.001)	(0.010)	(0.010)	(0.009)	(0.001)	(0.001)	(0.001)
Statement Related	−0.000	−0.000*	−0.000	−0.001	−0.001	−0.000	0.001**	0.001**	0.000
	(0.000)	(0.000)	(0.000)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.000)
FLS_Ratio	0.004*	0.005*	0.005**	0.005	0.006	0.020	0.001	0.000	−0.000
	(0.003)	(0.003)	(0.002)	(0.017)	(0.017)	(0.015)	(0.001)	(0.001)	(0.001)
Δ FDFD	−0.011	−0.008	0.000	0.088	0.040	0.000	−0.012	0.000	0.007
	(0.018)	(0.018)	(.)	(0.185)	(0.178)	(.)	(0.015)	(.)	(0.013)
MPU	−0.000			0.001	0.000		0.000	0.000	−0.000
	(0.000)			(0.000)	(0.000)		(0.000)	(.)	(0.000)
Predrift SPY	−0.001	−0.001	0.000						
	(0.005)	(0.006)	(.)						
Hawkish	−0.000	−0.000	−0.000	0.002	0.002	0.000	−0.000***	−0.000***	−0.000

	(0.000)	(0.000)	(0.000)	(0.003)	(0.003)	(0.003)	(0.000)	(0.000)	(0.000)
Public_Interest			0.000			0.000	0.001***	0.000	
			(.)			(.)	(0.000)	(.)	
Predrift VIX				−0.006	−0.006	0.000			
				(0.007)	(0.007)	(.)			
Predrift EUR							0.005	0.000	
							(0.005)	(.)	
Predrift JPY									0.010
									(0.011)
Chair FE	No	Yes	No	No	Yes	No	No	No	No
Meeting FE	No	No	Yes	No	No	Yes	No	Yes	No
r^2	0.016	0.290	0.049	0.049	0.065	0.316	0.044	0.235	0.034
N	1389.000	1389.000	1359.000	1359.000	1404.000	1359.000	1404.000	1404.000	1404.000

Table 5. Do investors react to negative facial expressions? This table presents coefficients from OLS regressions examining changes in stock volume (SPY), currency (EUR) tick count, and the VIX volatility index tick count in response to FOMC chairs' negative emotions and control variables. The analysis includes 1359 to 1404 observations at the minute level spanning 46 FOMC meetings chaired by Ben Bernanke (12), Janet Yellen (16), and Jerome Powell (18) from April 27th, 2011, to September 16th, 2020. Percent changes in SPY, VIX, EUR, JPY are measured over each minute and the absolute value is taken. Negative Facial Expressions represents the intensity of chairs' negative emotions averaged over the preceding three minutes relative to the average across all meetings under the chair. This is to control for the specific nature and disposition of each Fed Chair. Negative Sentiment measures the expressed tone based on FinBERT for each statement . Hawkishness measures the policy stance of chairs based on the keyword list in (Neuhierl and Webet, 2019) and spaCy LLM tokenization. Statement Related measures the frequency of statements in a time interval that are related to the FOMC Press Statement given at 2.00pm. All language parameters are averaged over each rolling minute. Predrift captures percent changes in the 30 minutes from 2.00pm to 2.30pm before the FOMC press conference for SPY, VIX, EUR, JPY respectively. %ΔFDFD denotes the change in Federal Funds Rate on the day of FOMC Press Conference. Standard errors, shown in parentheses, are heteroskedasticity-robust. Variable definitions are detailed in Table 1.

	(1) SPY Vol	(2) SPY Vol	(3) VIX Tick	(4) VIX Tick	(5) EUR Tick	(6) EUR Tick	(7) EUR Tick
Negative Facial	−0.035	−0.059***	−0.453***	−0.002	0.386***	−0.501***	−0.068
	(0.027)	(0.023)	(0.079)	(0.004)	(0.067)	(0.069)	(0.078)
Negative Sentiment	0.019	0.022	0.004	−0.003	0.168***	0.141***	0.095***
	(0.013)	(0.017)	(0.044)	(0.003)	(0.055)	(0.046)	(0.033)
Statement Related	−0.004*	−0.003	−0.005	0.001	0.028***	0.022***	0.023***
	(0.002)	(0.002)	(0.010)	(0.001)	(0.009)	(0.007)	(0.005)
Hawkish	−0.002	−0.001	0.002	0.000	−0.000	−0.003	−0.007*
	(0.002)	(0.002)	(0.006)	(0.000)	(0.005)	(0.005)	(0.004)
Δ FDFD	0.000	−0.463***	−1.475***	0.000	−1.050***	−1.404***	0.000
	(.)	(0.082)	(0.172)	(.)	(0.213)	(0.180)	(.)
Predrift SPY	0.000	−0.052					
	(.)	(0.041)					
Predrift VIX			−0.105***	0.000			
			(0.009)	(.)			
Predrift EUR					0.697***	0.432***	0.000
					(0.141)	(0.121)	(.)

Chair FE	No	Yes	Yes	No	No	No	Yes
Meeting FE	No	No	No	Yes	No	No	Yes
r^2	0.399	0.122	0.461	0.998	0.069	0.375	0.691
N	1389.000	1389.000	1359.000	1359.000	1404.000	1404.000	1404.000

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

Table 6. Does the Fed Chair consciously control his/her facial expression? This table presents coefficients from OLS regressions examining changes in stock (SPY), currency (EUR), (JPY), and the VIX volatility index in response to FOMC chairs' negative facial expressions and control variables. The analysis includes 1404 observations at the minute level spanning 46 FOMC meetings chaired by Ben Bernanke (12), Janet Yellen (16), and Jerome Powell (18) from April 27th, 2011, to September 16th, 2020. Percent changes in SPY, VIX, EUR are measured over each minute and the absolute value is taken. The dependent variables are the facial expressions. Negative Facial Expressions represents the intensity of chairs' emotions related to the particular emotion averaged over the preceding three minutes relative to the average across all meetings under the chair. This is to control for the specific nature and disposition of each Fed Chair. Negative Sentiment measures the expressed tone based on FinBERT for each statement . Hawkishness measures the policy stance of chairs based on the keyword list in (Neuhierl and Webet, 2019) and spaCy LLM tokenization. Statement Related measures the frequency of statements in a time interval that are related to the FOMC Press Statement given at 2.00pm. All language parameters are averaged over each rolling minute. Predrift captures percent changes in the 30 minutes from 2.00pm to 2.30pm before the FOMC press conference for SPY, VIX, EUR respectively. Standard errors, shown in parentheses, are clustered at the meeting level. Variable definitions are detailed in Table 1.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Negative	Negative	Negative	Neutral	Neutral	Happy	Sad
Conference Count	0.031***	0.035***	0.000	−0.059***	0.000	−0.023	−0.027
	(0.003)	(0.006)	(.)	(0.013)	(.)	(0.021)	(0.016)
Negative Sentiment	−0.011	−0.016	−0.008	0.030	0.039*	0.037	−0.024
	(0.017)	(0.012)	(0.011)	(0.026)	(0.020)	(0.047)	(0.025)
Statement Related	−0.001	−0.003	−0.000	0.004	0.003	−0.010	0.006
	(0.003)	(0.003)	(0.003)	(0.006)	(0.005)	(0.007)	(0.006)
Δ FDFD	−0.162*	−0.199	0.000	−0.297	0.000	0.175	0.818*
	(0.088)	(0.282)	(.)	(0.340)	(.)	(0.835)	(0.441)
MPU	−0.001***	−0.001	0.000	0.003**	0.000	−0.003**	0.002
	(0.000)	(0.001)	(.)	(0.001)	(.)	(0.001)	(0.001)
Chair FE	No	Yes	No	Yes	No	Yes	Yes
Meeting FE	No	No	Yes	No	Yes	No	No
r2	0.089	0.329	0.558	0.194	0.456	0.064	0.154
N	1404.000	1404.000	1404.000	1404.000	1404.000	1404.000	1404.000

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

Table 7. Do investors react to other facial expressions that convey transparency? This table presents coefficients from OLS regressions examining changes in stock (SPY), currency (EUR), (JPY), and the VIX volatility index in response to FOMC chairs' happy facial expressions, transparent facial expressions, sad facial expressions, neutral facial expressions and control variables. The analysis includes 1359 to 1404 observations at the minute level spanning 46 FOMC meetings chaired by Ben Bernanke (12), Janet Yellen (16), and Jerome Powell (18) from April 27th, 2011, to September 16th, 2020. Percent changes in SPY, VIX, EUR are measured over each minute and the absolute value is taken. Transparent Facial Expression represents both neutral and happy. The Facial Expressions represents the intensity of chairs' emotions related to the particular emotion averaged over the preceding three minutes relative to the average across all meetings under the chair. This is to control for the specific nature and disposition of each Fed Chair. Negative Sentiment measures the expressed tone based on FinBERT for each statement . Hawkishness measures the policy stance of chairs based on the keyword list in (Neuhierl and Webet, 2019) and spaCy LLM tokenization. Statement Related measures the frequency of statements in a time interval that are related to the FOMC Press Statement given at 2.00pm. All language parameters are averaged over each rolling minute. Predrift captures percent changes in the 30 minutes from 2.00pm to 2.30pm before the FOMC press conference for SPY, VIX, EUR respectively. %ΔFDFD denotes the change in Federal Funds Rate on the day of FOMC Press Conference. Standard errors, shown in parentheses, are heteroskedasticity-robust. Variable definitions are detailed in Table 1.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	%Δ SPY	%Δ SPY	%Δ SPY	%Δ SPY	%Δ VIX	%Δ VIX	%Δ VIX	%Δ EUR	%Δ EUR
Happy Facial	−0.001*								−0.002***
	(0.001)								(0.001)
Negative Sentiment	0.003**	0.003*	0.003**	0.003**	0.011	0.008	0.013	0.001	0.001
	(0.002)	(0.002)	(0.001)	(0.002)	(0.014)	(0.014)	(0.012)	(0.001)	(0.001)
Hawkish	−0.000	−0.000	−0.000	−0.000	0.002	0.001	0.000	−0.000***	−0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.003)	(0.003)	(0.003)	(0.000)	(0.000)
Statement Related	−0.000**	−0.000**	−0.000	−0.000**	−0.001	−0.001	−0.000	0.001**	0.001**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
%Δ FDFD	−0.012*	−0.011	0.000	−0.017***	0.059	0.026	0.000		
	(0.007)	(0.007)	(.)	(0.007)	(0.066)	(0.065)	(.)		
MPUUS_MPU	−0.000***	−0.000***	0.000	−0.000***	0.000	−0.000	0.000		
	(0.000)	(0.000)	(.)	(0.000)	(0.000)	(0.000)	(.)		
Predrift SPY	−0.002	−0.002	0.000	−0.002					
	(0.002)	(0.002)	(.)	(0.002)					
Neutral Facial		0.003*			0.012				

		(0.002)			(0.013)				
Transparent Facial			−0.006***			−0.068***		−0.002*	
			(0.002)			(0.014)		(0.001)	
Sad Facial				0.007***			0.014		
				(0.002)			(0.023)		
Predrift VIX					−0.006**	−0.006**	0.000		
					(0.003)	(0.003)	(.)		
Predrift EUR								0.003	0.004
								(0.003)	(0.003)
Chair FE	Yes	Yes	No	Yes	Yes	Yes	No	Yes	Yes
Meeting FE	No	No	Yes	No	No	No	Yes	No	No
r2	0.031	0.032	0.285	0.042	0.042	0.042	0.315	0.315	0.064
N	1389.000	1389.000	1389.000	1389.000	1359.000	1359.000	1359.000	1359.000	1404.000

Table 8. Do investors react to facial expressions that contrast with word sentiment? This table presents coefficients from OLS regressions examining changes in stock (SPY), currency (EUR), (JPY), and the VIX volatility index in response to FOMC chairs' negative facial expressions and control variables. The analysis includes 1359 to 1420 observations at the minute level spanning 46 FOMC meetings chaired by Ben Bernanke (12), Janet Yellen (16), and Jerome Powell (18) from April 27th, 2011, to September 16th, 2020. Percent changes in SPY, VIX, EUR are measured over each minute and the absolute value is taken.Negative Facial Expressions represents the intensity of chairs' emotions related to the particular emotion averaged over the preceding three minutes relative to the average across all meetings under the chair. This is to control for the specific nature and disposition of each Fed Chair. Negative Sentiment measures the expressed tone based on FinBERT for each statement . Hawkishness measures the policy stance of chairs based on the keyword list in (Neuhierl and Webet, 2019) and spaCy LLM tokenization. Statement Related measures the frequency of statements in a time interval that are related to the FOMC Press Statement given at 2.00pm. All language parameters are averaged over each rolling minute. Predrift captures percent changes in the 30 minutes from 2.00pm to 2.30pm before the FOMC press conference for SPY, VIX, EUR respectively. Standard errors, shown in parentheses, are clustered at the meeting level. Variable definitions are detailed in Table 1.

	(1)	(2)	(3)	(4)	(5)
	%Δ SPY	%Δ SPY	%Δ VIX	%Δ VIX	%Δ EUR
Negative Facial	−0.007	−0.006**	−0.085**	−0.054*	−0.003
	(0.006)	(0.003)	(0.039)	(0.033)	(0.003)
Negative Sentiment	0.004	0.002	0.005	0.014	−0.001
	(0.004)	(0.003)	(0.029)	(0.033)	(0.003)
Negative Facial * Negative_Sent	−0.001	0.000	0.003	−0.002	0.002
	(0.003)	(0.002)	(0.026)	(0.026)	(0.002)
Statement Related	−0.000		−0.001		0.001*
	(0.000)		(0.002)		(0.001)
%Δ FDFD	−0.009		0.035		−0.017
	(0.018)		(0.184)		(0.012)
Predrift SPY	−0.001	0.000			
	(0.006)	(.)			
Predrift VIX			−0.006	0.000	
			(0.007)	(.)	
Predrift EUR					0.002
					(0.005)

Chair FE	No	Yes	Yes	No	Yes
Meeting FE	No	Yes	No	Yes	No
r2	0.022	0.279	0.049	0.315	0.065
N	1389.000	1450.000	1359.000	1420.000	1404.000

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

Table 9. How do investors interpret negative facial expressions? This table presents coefficients from OLS regressions examining changes in stock (SPY) and the VIX volatility index in response to FOMC chairs' negative facial expressions and control variables. The analysis includes 1404 observations at the minute level spanning 46 FOMC meetings chaired by Ben Bernanke (12), Janet Yellen (16), and Jerome Powell (18) from April 27th, 2011, to September 16th, 2020. Percent changes in SPY, VIX, EUR are measured over each minute and the absolute value is taken. Negative Facial Expressions represents the intensity of chairs' emotions related to the particular emotion averaged over the preceding three minutes relative to the average across all meetings under the chair. This is to control for the specific nature and disposition of each Fed Chair. Negative Sentiment measures the expressed tone based on FinBERT for each statement . Hawkishness measures the policy stance of chairs based on the keyword list in (Neuhierl and Webet, 2019) and spaCy LLM tokenization. Statement Related measures the frequency of statements in a time interval that are related to the FOMC Press Statement given at 2.00pm. All language parameters are averaged over each rolling minute. Cfquart is an indicator, factor variable of which quartile of a Fed Chair's career is that conference in. Congress30 and Congress10 represent whether a FOMC Press Conference was held within 30 days or 10 days after a congressional testimony respective. Predrift captures percent changes in the 30 minutes from 2.00pm to 2.30pm before the FOMC press conference for SPY, VIX respectively. Standard errors, shown in parentheses, are heteroskedasticity-robust. Variable definitions are detailed in Table 1.

	(1) %Δ SPY	(2) %Δ SPY	(3) %Δ SPY	(4) %Δ SPY	(5) %Δ SPY	(6) %Δ VIX	(7) %Δ VIX	(8) %Δ VIX	(9) %Δ VIX
Negative Facial	−0.005	−0.000	−0.004	−0.009***	−0.008***	−0.015	0.036	−0.029	−0.079***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.028)	(0.026)	(0.019)	(0.021)
Negative Facial i/r Conference Count	−0.000	−0.001***				−0.005***			
	(0.000)	(0.000)				(0.002)			
Negative Sentiment	0.003*	0.003*	0.003*	0.003*	0.003*	0.009	0.009	0.009	0.009
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.014)	(0.014)	(0.014)	(0.014)
Statement Related	−0.000**	−0.000**	−0.000**	−0.000**	−0.000**	−0.001	−0.000	−0.000	−0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.002)	(0.002)	(0.002)	(0.002)
%Δ FDFD	−0.010	−0.006	−0.010	−0.012*	−0.008	0.066	0.122*	0.148*	0.018
	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)	(0.068)	(0.070)	(0.078)	(0.067)
MPUUS_MPU	−0.000		−0.000***	−0.000***	−0.000	0.000	0.001***	0.001***	0.000
	(0.000)		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Predrift SPY	−0.001	0.000	−0.002	−0.002	−0.000				
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)				
Negative * cfquart			−0.002**				−0.023***		

			(0.001)				(0.007)		
Negative * congre 30				0.001				0.038**	
				(0.002)				(0.019)	
Negative * congre 10					0.007*				−0.035
					(0.004)				(0.023)
Predrift VIX						−0.006**	−0.005**	−0.007***	−0.005**
						(0.003)	(0.002)	(0.002)	(0.003)
Chair FE	No	Yes	Yes	Yes	No	Yes	No	No	Yes
r2	0.012	0.028	0.039	0.037	0.014	0.054	0.017	0.013	0.050
N	1389.000	1389.000	1389.000	1389.000	1389.000	1359.000	1359.000	1359.000	1359.000

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

Appendix

A1. FOMC Press Conference and Introductory Statement Recordings List

Panel A: List of Publicly Available Recordings of FOMC Press Conference		
4/27/2011	6/22/2011	11/2/2011
1/25/2012	4/25/2012	6/20/2012
9/13/2012	12/12/2012	3/20/2013
6/19/2013	9/18/2013	12/18/2013
3/19/2014	6/18/2014	9/17/2014
12/17/2014	3/18/2015	6/17/2015
9/17/2015	12/16/2015	3/16/2016
6/15/2016	9/21/2016	12/14/2016
3/15/2017	6/14/2017	9/20/2017
12/13/2017	3/21/2018	6/13/2018
9/26/2018	12/19/2018	1/30/2019
3/20/2019	5/1/2019	6/19/2019
7/31/2019	9/18/2019	10/30/2019
12/11/2019	1/29/2020	4/29/2020
6/10/2020	7/29/2020	9/16/2020
11/5/2020	12/16/2020	

Appendix A1: List of 47 dates of FOMC Press Conference video recordings publicly available and used in this paper. All press conferences start at 2.30pm EST and we verify the timestamps through the livestream on YouTube. For videos with introductory statements (a feature introduced by Powell), the press conference continues immediately after the statement. I left out March 3rd 2020 due to incompleteness of the data. Press conferences from 29th April 2020 onwards are done via zoom due to Covid restrictions.

A2. Video Background Analysis

Due to the large number of screenshots taken at fixed 2-second intervals, I separate out those where reports are asking questions from the Fed Chair using a background analyzer. For instance, on 5th November 2020, screenshot_0827.jpg shows Powell speaking but screenshot_0828 shows the camera panning out to the zoom session. The program correctly identifies using cosine similarity test that screenshot_0828 is not of the Fed Chair, thus, I remove these screenshots from the analysis.

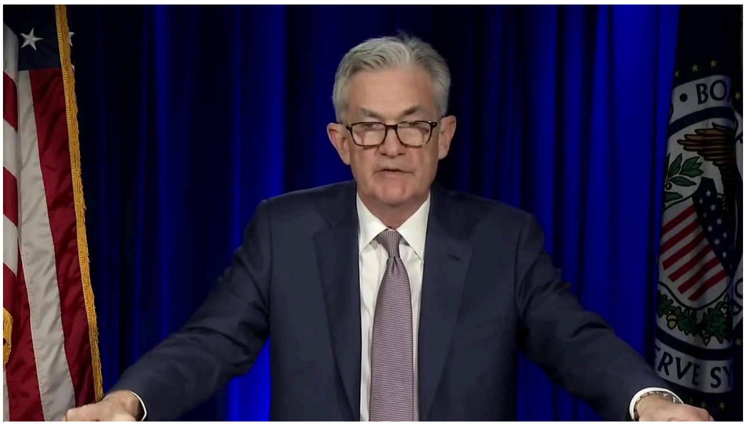

Panel A. Fed Chair speaking in *screenshot_0827.jpg*, which the program identifies correctly as an image of Powell speaking, thus, I include it in the analysis.

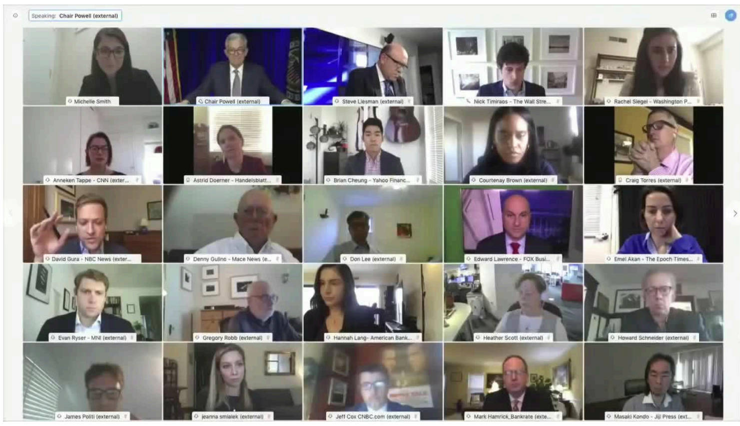

Panel B. Reporters asking questions in *screenshot_0827.jpg*, which the program identifies correctly as not an image of Powell speaking. I remove it from the analysis.

Figure A1. Comparison of separate screenshots of Fed Chair speaking versus reporters asking questions.

A3. NLP Keywords

3. To check for hawkish sentiments, I employ the keyword search by Neuhierl and Weber (2019).

Dovish	Hawkish
anchor inflation expectations	aggregate demand higher
anchored inflation expectations	asset prices increase
boost aggregate demand	asset prices rise
boost economic activity	business investment increased
cut federal funds rate	declines unemployment rate
cut interest rates	declining unemployment rate
cuts federal funds rate	drop unemployment rate
cutting federal funds rate	economic activity increased
declines asset prices	economic outlook increased
declines crude oil	employment increased
declines economic activity	energy prices rise
declines employment	exchange rates lower
declines energy prices	gradual increases federal funds rate
declines house prices	gross domestic product rising
declines labor force participation	growing current account deficit
declining house prices	higher asset prices
declining interest rates	higher employment
downward pressure asset prices	higher energy prices
downward pressure house prices	higher federal funds rate
downward pressure interest rates	higher house prices
drop crude oil	higher inflation expectations
drop house prices	higher interest rates
eased stance monetary policy	higher productivity growth
easing monetary policy	higher unit labor costs
employment declined	house prices increase
employment fallen	house prices increased
employment fell	house prices rise
employment stable	house prices rising
federal funds rate lower	increase asset prices
firmly anchored inflation expectations	increase core inflation
house prices declined	increase current account surpluses

house prices fallen	increase economic activity
house prices fell	increase employment
increase aggregate demand	increase energy prices
increase current account deficit	increase federal funds rate
increase labor productivity	increase house prices
increase unemployment rate	increase inflation expectations
increases productivity growth	increase interest rates
increases labor productivity	increase productivity growth
increases productivity growth	increase resource utilization
inflation expectations anchored	increase target federal funds
inflation expectations declined	increase unit labor costs
inflation expectations firmly anchored	increased economic activity
inflation expectations remained stable	increased employment
inflation expectations stable	increased labor force participation
inflation expectations well anchored	increases aggregate demand
interest rates declined	increases asset prices
interest rates drop	increases business investment
interest rates easing	increases crude oil
interest rates lower	increases employment
interest rates lowering	increases energy prices
interest rates remain	increases federal funds rate
keeping interest rates	increases house prices
keeping monetary policy	increases inflation expectations
labor productivity increased	increases interest rates
lower energy prices	increases output gap
lower federal funds rate	increases unit labor costs
lower house prices	inflation expectations increased
lower inflation expectations	interest rates higher
lower interest rates	interest rates increase
lower level real oil prices	interest rates increased
lower potential output	interest rates might rise
lowered federal funds rate	interest rates raise
lowering federal funds rate	interest rates raised
lowering interest rates	interest rates rise
monetary policy easing	interest rates rising

nonaccelerating inflation rate	lower current account deficit
productivity growth increased	lower productivity growth
productivity growth increases	lower unemployment rate
raise aggregate demand	monetary policy tightening
rapid productivity gains	personal saving rate fallen
reduce federal funds rate	raise federal funds rate
reduce interest rates	raise interest rates
reduce unemployment rate	raised interest rates
reduced economic activity	raising asset prices
reduced federal funds rate	raising federal funds rate
reduced interest rates	raising interest rates
reducing federal funds rate	rapid productivity growth
reducing interest rates	reduce current account deficit
reduction aggregate demand	reductions unemployment rate
reduction federal funds rate	resource utilization increased
reduction inflation expectations	rise asset prices
reduction interest rates	rise core inflation
reductions federal funds rate	rise employment
reductions interest rates	rise energy prices
resource utilization subdued	rise federal funds rate
rise productivity growth	rise headline inflation
rise unemployment rate	rise house prices
rising current account deficit	rise inflation expectations
rising productivity growth	rise interest rates
risks economic activity	rise personal saving rate
risks economic outlook	rise unit labor costs
risks outlook economic activity	rising asset prices
stabilizing economic activity	rising employment
stabilizing employment	rising energy prices
stabilizing monetary policy	rising house prices
stable economic conditions	rising inflation expectations
stable inflation expectations	rising interest rates
stable inflation rate	risks long term inflation outlook
stable interest rates	sharp increases energy prices
stable monetary policy	sharp increases interest rates

stable prices moderate	sharp rise interest rates
subdued unit labor costs	tightening monetary policy
sustainable employment	unemployment rate declining
unemployment rate declined	unemployment rate fallen
unemployment rate rising	unemployment rate fell
upward pressure exchange rates	unemployment rate lower
well anchored inflation expectations	upward pressure core inflation
	upward pressure interest rates

2. To check for Statement related sentences, I compile a list of keywords after manually observing a representative sample of the transcripts.

Statement Related Terms		
FOMC statement	earlier statement	stated then
todays policy	indicated in the statement	early this afternoon
policy statement	released earlier today	conjunction with meeting
todays meeting	extensive discussions	what monetary policy might
turning to todays meeting	our economic outlook	our projection
we expect inflation	our views	our operations
our work	our statement	our measure
our strategy	contingent on projected	growth is expected
committees	committees belief	our response
our guess		

3. To check for forward-looking sentences (FLS), I use FinBERT FLS (Huang, Hui and Yi, 2023). The FLS extension identifies statements containing projections, anticipations, or references to future events or conditions. This method enhances the ability to control for statements that inherently signal expectations and policy forecasts, thus improving the precision of isolating nonverbal communication effects.

A4. Deepfake Video Creation

To create the deepfakes, I use DeepFaceLab, Nvidia RTX3000 series, build 11.20.2021 to train my models. I first identify the source and destination videos from PCA analysis to find the representative videos.

I use a computer specification of GPU - Nvidia GeForce RTX 3060, and a CPU - 12th Gen Intel i7-12700KF, 3610 Mhz. I use the default settings when training the model, except for specific instances where a different setting may outperform the default. For example, I use the *df-ud* model in training the SAEHD framework and use batch size of 4. More details can be found in Figure A2.

To corroborate the results, I use Faceswap.py. For the settings, I use the default. Color adjustment is set to "Match-Hist" for better blending between source and target faces. The "Extended" mask type is selected to include more facial features during merging. Output is set to use OpenCV, with the frame scale at 100%, keeping the original video resolution. Face processing options include a face scale of 0.0 and a reference threshold of 0.4 to control alignment quality. These settings aim to ensure realistic face swaps while preserving video clarity.

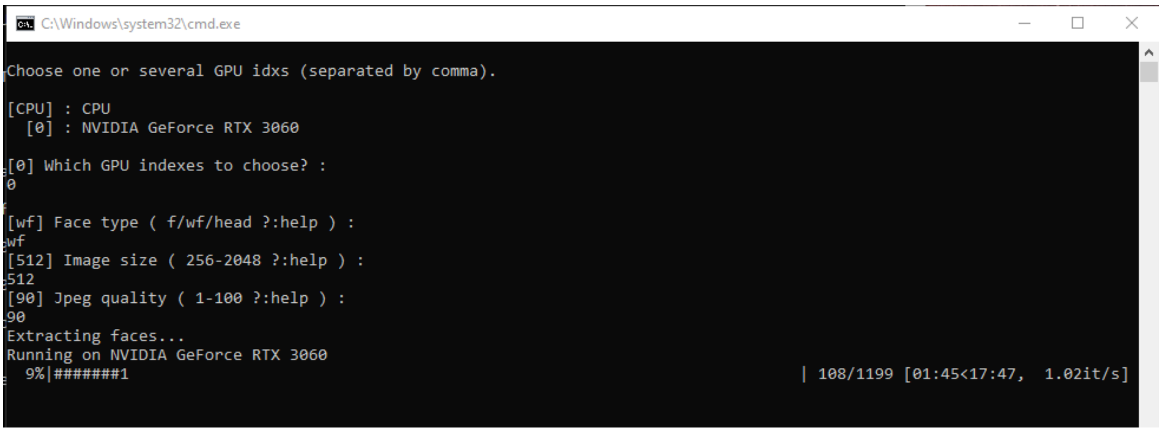

Panel A: Specification used to extract the faces from the videos

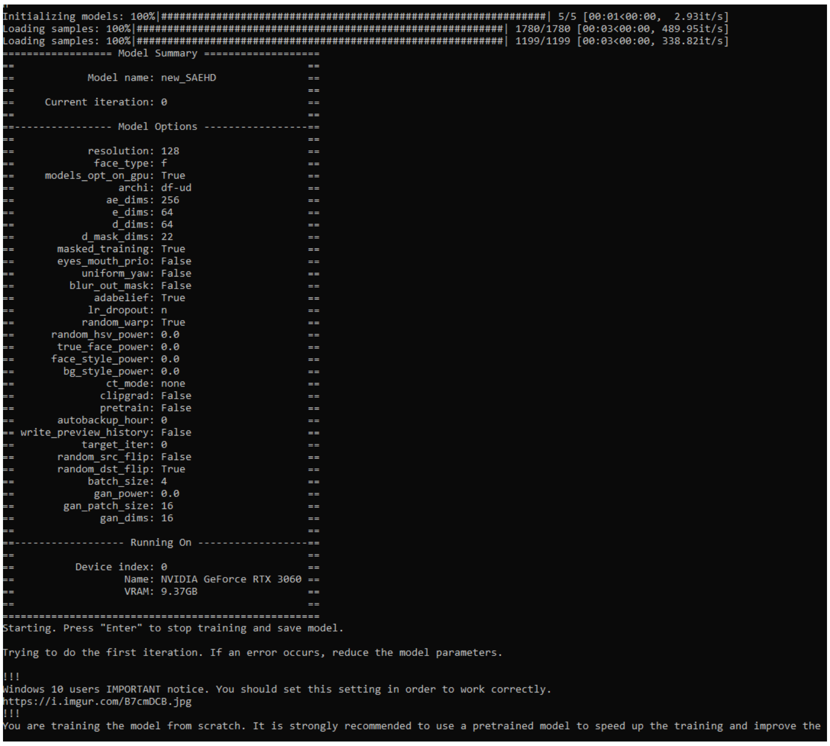

```
Initializing models: 100%|##################################################################| 5/5 [00:01<00:00,  2.93it/s]
Loading samples: 100%|##################################################################| 1780/1780 [00:03<00:00, 489.95it/s]
Loading samples: 100%|##################################################################| 1199/1199 [00:03<00:00, 338.82it/s]
================= Model Summary ==================
==                                              ==
==            Model name: new_SAEHD             ==
==                                              ==
==     Current iteration: 0                     ==
==                                              ==
==---------------- Model Options ---------------==
==                                              ==
==            resolution: 128                   ==
==             face_type: f                     ==
==     models_opt_on_gpu: True                  ==
==                 archi: df-ud                 ==
==               ae_dims: 256                   ==
==                e_dims: 64                    ==
==                d_dims: 64                    ==
==           d_mask_dims: 22                    ==
==       masked_training: True                  ==
==       eyes_mouth_prio: False                 ==
==           uniform_yaw: False                 ==
==         blur_out_mask: False                 ==
==             adabelief: True                  ==
==            lr_dropout: n                     ==
==           random_warp: True                  ==
==      random_hsv_power: 0.0                   ==
==       true_face_power: 0.0                   ==
==      face_style_power: 0.0                   ==
==        bg_style_power: 0.0                   ==
==               ct_mode: none                  ==
==              clipgrad: False                 ==
==              pretrain: False                 ==
==       autobackup_hour: 0                     ==
== write_preview_history: False                 ==
==           target_iter: 0                     ==
==       random_src_flip: False                 ==
==       random_dst_flip: True                  ==
==            batch_size: 4                     ==
==             gan_power: 0.0                   ==
==        gan_patch_size: 16                    ==
==              gan_dims: 16                    ==
==                                              ==
==------------------ Running On ----------------==
==                                              ==
==          Device index: 0                     ==
==                  Name: NVIDIA GeForce RTX 3060 ==
==                  VRAM: 9.37GB                ==
==                                              ==
==================================================
Starting. Press "Enter" to stop training and save model.

Trying to do the first iteration. If an error occurs, reduce the model parameters.

!!!
Windows 10 users IMPORTANT notice. You should set this setting in order to work correctly.
https://i.imgur.com/B7cmDCB.jpg
!!!
You are training the model from scratch. It is strongly recommended to use a pretrained model to speed up the training and improve the
```

Panel B. Comparison of a deepfake of Fed Chair Janet Yellen during FOMC press conference on September 21, 2016 using Fed Chair Powell and their facial analysis result

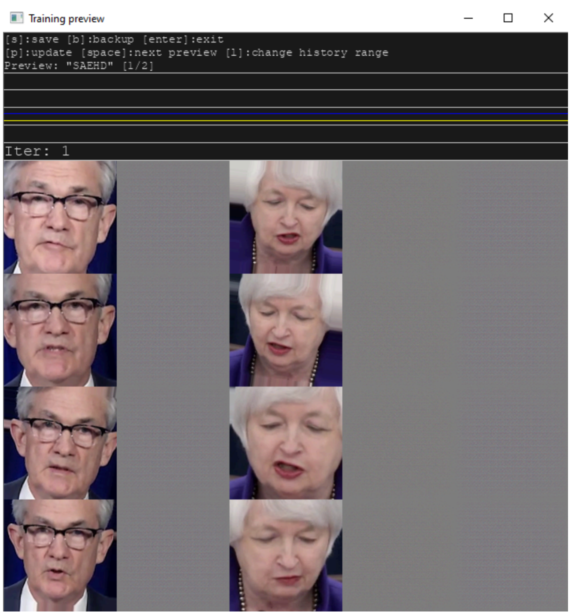

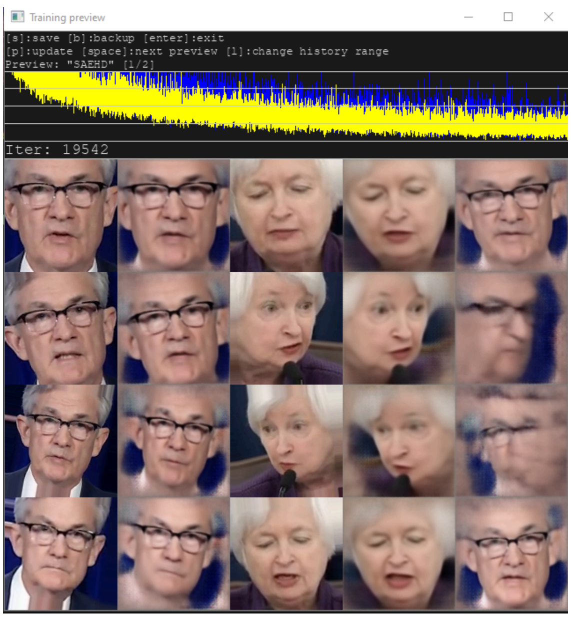

Panel C. Training of SAEHD mode, starting from iteration 1. I train to an average of a 100,000 iterations before constructing the deepfake videos.

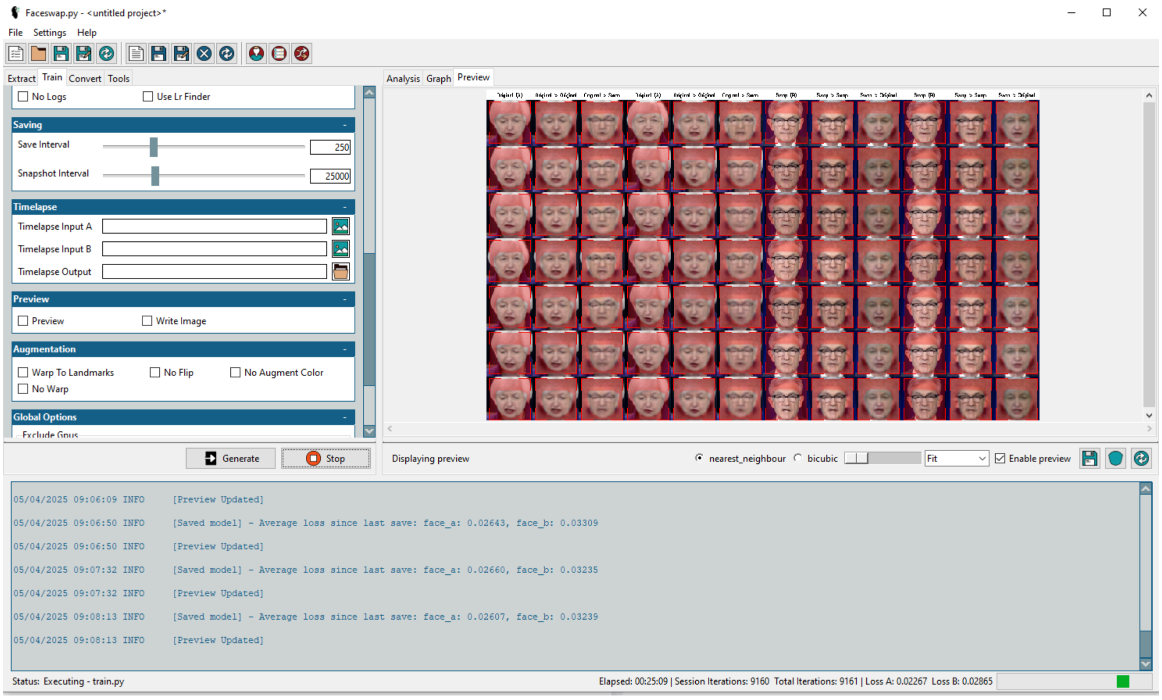

Panel D. Training using Faceswap.py. I train to an average of a 100,000 iterations before constructing the deepfake videos.

Figure A2. This figure documents the comprehensive pipeline for deepfake generation using DeepFaceLab (build 11.20.2021) on an Nvidia GeForce RTX 3060 GPU and 12th Gen Intel i7-12700KF CPU. Panel A details the face extraction process from source and destination videos, selected via PCA analysis. Panel B presents a comparative analysis of deepfakes involving Fed Chairs Janet Yellen and Jerome Powell, highlighting facial synthesis outcomes. Panel C outlines the SAEHD training configuration using the df-ud model architecture with a batch size of 4, converging at 100,000 iterations. Panel D illustrates the Faceswap.py training protocol, also conducted over 100,000 iterations. This visual sequence shows the technical and procedural steps underlying synthetic facial video construction.